



City Research Online

City St George's, University of London

Citation: Charness, G. & Sontuoso, A. (2023). The Doors of Perception: Theory and Evidence of Frame-Dependent Rationalizability. *American Economic Journal: Microeconomics*, 15(2), pp. 309-344. doi: 10.1257/mic.20210082

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/33727/>

Link to published version: <https://doi.org/10.1257/mic.20210082>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

The Doors of Perception:
Theory and Evidence of Frame-Dependent Rationalizability

Gary Charness *and* Alessandro Sontuoso

August 9, 2022

Abstract. We investigate how strategic behavior is affected by the set of notions (frames) used when thinking about the game. In our games, the action set consists of visual objects: each player must privately choose one, trying to match the counterpart's choice. We propose a model where different player-types are aware of different attributes of the action set (hence, different frames). One of the novelties is an epistemic structure that allows players to think about new frames, after initial *unawareness* of some attributes. To test the model, our experimental design brings about multiple frames by varying subjects' awareness of several attributes.

KEYWORDS: unawareness, awareness, interactive epistemology, rationalizability, coordination.

JEL Classification Numbers: C72, C91, D83.

Acknowledgments: We are grateful to Ned Augenblick, Alberto Cardaci, Vince Crawford, Alexander Funcke, Einav Hart, Rosemarie Nagel, David Rojo Arjona, Burkhard Schipper, Bob Sugden, Emanuel Vespa, and three anonymous referees for their helpful comments. Also, we thank Michael Cooper for running the experimental sessions and for programming the zTree code used in the experiment.

Contact: Gary Charness, *Department of Economics*, University of California, Santa Barbara, California, 93106, charness@econ.ucsb.edu; Alessandro Sontuoso, *Smith Institute for Political Economy and Philosophy*, Chapman University, One University Dr., Orange, California, 92866 and *Philosophy, Politics and Economics*, University of Pennsylvania, 249 S. 36th St., Philadelphia, Pennsylvania, 19104, sontuoso@sas.upenn.edu.

I. Introduction

Cognitive scientists define a “frame” as a bundle of information about the typical characteristics of a situation or problem. Frames are stored in individuals’ minds and provide default information with which to interpret and respond to events (Schank and Abelson, 1977). Similarly, artificial-intelligence pioneer Marvin Minsky (1975) codified frames so as to represent *knowledge structures* in the context of visual-reasoning and communication-processing problems: for instance, the correct interpretation of (and response to) a hand gesture depends on the agents’ understanding of the situation in which the gesture occurs. Indeed, in many everyday interactions individuals face an implicit coordination puzzle, whose solution depends on the set of notions (i.e., the frames) they take into account when considering it. Many such problems might seem trivial, since we often interact with people with whom we share a similar frame, to the point that one does not even notice that there was a coordination puzzle in the first place. However, miscoordination may become evident when one interacts with people who see the problem through a different lens (think of the awkward confusion about greeting styles we sometimes experience in social settings).¹

The cognitive processes underlying everyday coordination problems might be thought of in this way: first, one “mentally frames” (i.e., describes) the problem on the basis of its perceived characteristics; then, one follows the solution that is naturally associated (i.e., comes to mind easily) with the relevant frame. A key determinant of an individual’s behavior is therefore her perception of the problem’s characteristics. Accordingly, in this paper we shall investigate – theoretically and experimentally – how choice behavior in abstract coordination games is affected by the individual’s perception and beliefs about others’ perception.

In formal accounts of coordination problems (i.e., symmetric, simultaneous-move games with multiple pure-strategy Nash equilibria),² the issue of singling out an optimal course of action has been a long-standing consideration, due to the non-uniqueness of possible solutions. Relatedly,

¹ In the economics literature, the theory of the firm views coordination problems as one of two key organization hurdles (the other being the much more studied “agency problem”; Milgrom and Roberts, 1992). It has been suggested that the coordination problem of organizations is inherently due to people’s cognitive limitations, in that individuals often lack a common understanding of the tasks they must integrate and coordinate upon (e.g., Heath and Staudenmayer, 2000). This line of research implies that individuals come to develop a different understanding of their tasks as a result of a different focus or background; different viewpoints may in turn entail different solutions to an identical or similar task (Arrow, 1974; Kreps, 1990; Okhuysen and Bechky, 2009; Kets and Sandroni, 2021).

² The class of coordination problems contains any situation in which there exist multiple ways players may “match” their behavior for mutual benefit. This class contains a broad and diverse array of interactions, including games with slightly conflicting interests and games *with* or *without* Pareto-rankable equilibria. In this paper, we consider “pure” coordination games (i.e., if players choose the same action, they each receive an identical positive payoff; otherwise, they each receive nothing), with such games featuring non-Pareto-rankable equilibria.

Thomas Schelling (1960) noted that the use of “situational cues” (i.e., characteristics or attributes of the problem) could help individuals converge on one solution; in fact, mutually-recognizable attributes often induce a frame in such a way to make a specific course of action come to mind readily. Here, we contribute to the analysis of frames by formalizing assumptions about attribute awareness, rationality and beliefs, and then by studying their behavioral implications in coordination problems. Below is an example.

Consider a “matching game” where two players are presented with a set of options, from which each player must privately choose one with the goal of matching the counterpart’s choice. In this case, a frame may be viewed as a player’s description of the options, based on the attributes she perceives and thinks about.³ For a schematic illustration, let’s suppose that player i ’s action set in a matching game consists of three visual objects denoted $\{a_1, a_2, a_3\}$, with options 1, 2, and 3 respectively representing a *cyan triangle*, a *cyan diamond*, and a *lavender triangle*. Note that the conventional way of defining a game does not permit a qualitative characterization of the options to enter the formal description of the game. Still, as was first suggested by philosopher David Gauthier (1975), accounting for the characteristics of the options implies “restructuring” the action set (e.g., someone who thinks about the objects’ colors would *distinguish actions* according to their colors).

Formally, the act of distinguishing between colored objects may be represented as the case where player i partitions the action set so that each of the cells corresponds to an instance of the color attribute; that is, $\{\{a_1, a_2\}, \{a_3\}\}$. By contrast, a player who thinks about the different shapes would partition the action set so that each cell corresponds to an instance of the shape attribute; i.e., $\{\{a_1, a_3\}, \{a_2\}\}$. Generalizing – for any attribute k – we say that k induces a frame for player i , if i partitions the action set so that each cell corresponds to an instance of attribute k : when that occurs, we say that i perceives and thinks about (“is aware of”) attribute k . (For example, the partition $\{\{a_2, a_3\}, \{a_1\}\}$ indicates that player i pays attention to some attribute with respect to which option 1 happens to be distinct.) Intuitively, players whose frames comprise several attributes will be aware of multiple ways to partition the action set.

Below we propose and test a model articulating what it entails for players to be *aware of* alternative frames. We shall see that the introduction of frame-dependent “awareness types”

³ For a real-world example, suppose that you are attending a conference out of town, and had previously agreed to meet a friend for a drink at the hotel’s bar at 9 pm. Once in the lobby, you realize that the hotel actually has three bars and your phone is out of battery. Under the assumption that each of you is indifferent about the meeting place and prefers having a drink together (rather than a drink alone), then you have a coordination problem where your choice behavior may ultimately depend on your perception of the bars’ salient characteristics.

restricts players' beliefs in a natural way, allowing us to rationalize why differences in players' frames may lead to differences in game play. Specifically, our belief restrictions imply that cells that contain fewer elements stand out, and thus are more attractive. So, in the example above, a player who thinks solely about attribute k (or else believes others do so) will end up choosing the option that is unique with respect to attribute k (i.e., the k "oddity").

As will be clear, our model integrates streams of research such as the analysis of "labelings" (e.g., Bacharach, 1993; Janssen, 2001) and the study of "unawareness", that is, the case where one does not know something *and* does not know that one does not know it (Dekel, Lipman, and Rustichini, 1998a; Heifetz, Meier, and Schipper, 2006).⁴ Bacharach's work was the first to explicitly define attribute-dependent labels so as to represent players' frames (see also Sugden, 1995; a related approach for the case of repeated games was pioneered by Crawford and Haller, 1990, where previous play would implicitly label actions in such a way as to generate a distinct option, e.g., "do what you did last time" vs. "do something else"). In particular, Bacharach's work was seminal in introducing a non-standard information structure, whereby player i does not know the list of all possible opponent types; rather, i only considers (and best-responds to) those types that think about attribute combinations i herself thinks about.⁵

While greatly innovative, the predictive scope of the labeling literature is somewhat limited by its solution concepts' reliance on strong assumptions (e.g., players' beliefs about the assignment of labels are consistent with an exogenous probability distribution, players use a set of principles for selecting among equilibria, etc.); thus, this literature is not well suited as a predictive tool in games where one has no experience about the others' perceptual limitations, or in cases where perception might change. To sidestep these drawbacks, we propose a new solution concept. In doing so, we

⁴ For a Level- k version of Bacharach's theory, see Bacharach and Stahl (2000). For early work on unawareness, see Fagin and Halpern (1987) and Modica and Rustichini (1994, 1999). In this connection, it is useful to see how unawareness differs from uncertainty: under *uncertainty*, one conceives of the space of all relevant contingencies (say, ω' , ω'' , ω'''), but does not know which has occurred and thus holds a probabilistic belief about each of them; under *unawareness*, instead, one does not think of all relevant contingencies (e.g., if one is unaware of the possibility that ω''' may occur, one will not hold any belief at all about ω'''). Dekel et al. (1998a) showed that unawareness cannot be accounted for under the assumptions underlying standard models of knowledge (in particular, "negative introspection" implies that at any state where an agent does not know some event, she *knows* that she *does not know* it, which rules out unawareness; see also Chen, Ely, and Luo, 2012, and for a survey, Schipper, 2014). To circumvent the Dekel et al. impossibility result, Heifetz et al. (2006) utilize a system of *multiple* state spaces, where different player types are aware of different spaces.

⁵ For example, suppose that the action set consists of marbles that vary in size, brilliance, and material: in Bacharach's (1993) model if a player ignores any such attributes, say, the marbles' size, then she will *not* hold any belief about the eventuality of facing a size-perceiving opponent. So, i will *not* best-respond to types that think about attributes i herself does not think about. Notably, Bacharach posited that i 's beliefs about the distribution of opponent types are consistent with the true exogenous distribution of types, or else consistent with a truncated and rescaled distribution if i ignores an attribute.

build on Heifetz et al.’s (2006, 2013a) account of unawareness, and explicitly incorporate an epistemic structure into the study of frames (i.e., unlike the Bacharach model and related work, we formally define a system of *multiple state spaces* to represent what individuals consider in regard to alternative frames and, hence, in regard to opponent types). Moreover, unlike the Heifetz et al. approach, we define a pair of *frame-dependent restrictions* on beliefs so as to characterize how each type may behave in a game. This leads to a novel solution concept (“frame-dependent rationalizability”), which we use to derive experimental predictions and provide a test of competing explanations for the role of frames in coordination games.⁶

To that end, we present an experiment involving ten (one-shot) matching games, played by pairs of subjects without feedback. Our between-subjects design brings about multiple frames by varying subjects’ awareness of several attributes. In examining the impact of attribute awareness on choice, we contrast our model’s predictions against a set of null hypotheses based on the “standard” single-state-space Bayesian paradigm (i.e., incomplete-information games with no unawareness).

In the *Baseline* treatment, each participant is shown six unlabeled objects (i.e., colored geometric shapes) on her screen. At the beginning of each game, the computer program selects three of those objects (the same three for each member of a pair, while the rest of the objects disappear from the screen). Participants are eventually prompted to choose one of the three available objects, with the goal of matching the counterpart’s choice; this entire set-up is common information. Each of the ten games presents a different three-object selection, and thus differs from the other games in the characteristics of the available action set.

The *All-Aware* treatment is the same as the Baseline, except that we hint at several attributes at once, in such a way that subjects are privately made aware of multiple attributes. To do so, we ask them how likely they think it would be for Baseline participants to notice each of the three attributes of the objects (i.e., color, shape, and order of display).⁷ These questions may be viewed as tautologies, as in “ E is the case or $\neg E$ is the case”, where E and $\neg E$ respectively represent the event such that others *have* and *have not* noticed an attribute. Note that a question involving a

⁶ Our solution concept is defined by an iterative deletion procedure, whereby each player i best-responds to her (subjective) beliefs about surviving type-action pairs of the opponents of which i is aware (beliefs about opponent types of which i is unaware remain undefined). If a player becomes aware of any additional frames, her optimization problem is updated so as to account for the eventuality of facing additional opponent types.

⁷ Specifically, we asked the following questions (in random order across games): “How likely do you think it is that the other participants have noticed the different colors of the objects? ... the different shapes? ... the order in which the objects have been drawn by the computer program?” To prevent these questions from generating common awareness of the frames, All-Aware participants are informed that the subjects with whom they are paired (i.e., Baseline participants) would *not* be asked such questions, but they would otherwise play the same game.

tautological clause conveys no information about whether E or $\neg E$ is true. Further, if one had been aware of E and $\neg E$ in the first place (i.e., if one had already paid attention to that attribute and thus considered the possibility that others might do so too), then any such question would not alter one's view of the game. By contrast, if one had *not* been aware of E and $\neg E$ in the first place, then the question itself would automatically generate awareness of those events. Indeed, a key goal of the All-Aware treatment is to ensure that its participants think about all the objects' attributes.

The individual-level data from the All-Aware treatment allow us to test if subjects optimize a frame-dependent expected utility (per our proposed notion of rationalizability). That is, we predict that each All-Aware participant chooses the oddity associated with the frame that, according to her own beliefs, is most-likely considered by her (Baseline) counterpart. Relatedly, we stress that the All-Aware treatment manipulation hints at several attributes at once, without directing subjects' attention to any one frame in particular, thereby minimizing any implicit demand effects. The data support our predictions, indicating that one's behavior depends on one's beliefs as to whether the counterpart has noticed (i.e., is aware of) an attribute or not. We will see that such intuitive results cannot be captured by "standard" Bayesian models (i.e., incomplete-information games with a single state space), which rule out the possibility that players may be unaware of anything relevant. Also, we will see that the data from the All-Aware treatment reveal substantial diversity in individuals' beliefs about others' awareness of the attributes (contradicting an implication of early models with non-standard information structures such as Bacharach, 1993).

Next, the Baseline vs. All-Aware treatment comparison permits us to test for heterogeneity in individuals' *awareness* of the attributes (as opposed to heterogeneity in individuals' *beliefs about others' awareness* of the attributes, which we discussed above). As a benchmark, we note that the assumptions underlying standard games with incomplete information (i.e., no unawareness) here imply no significant differences across treatments. In fact, if it were true that subjects knew all the possible frames/types in the first place (as is implied by the standard Bayesian paradigm), then participants would learn nothing from the All-Aware questions above (see footnote 7); also, since those questions in no way make an attribute more salient than another, they should *not* significantly change the All-Aware participants' choices. Yet, contrary to what is implied by previous models, our data reveal that the All-Aware manipulation does affect average game play. By comparing choice distributions across treatments, we conclude that the All-Aware manipulation makes some participants think about attributes to which they would otherwise not have paid attention (namely, order frames and, to a lesser extent, shape frames). The data further indicate that the All-Aware

manipulation leads to a decrease in coordination rates, as is predicted by our model: remarkably, this means that an increase in attribute awareness can hurt successful coordination, *ceteris paribus*.

To put the results in context, a few comments are in order. First, we stress that consistent with the epistemic literature, here being aware of an event (e.g., a frame) means that the event is being “taken into account” when one makes a choice (Modica and Rustichini, 1999, p. 274); that is, being aware corresponds to “thinking of” or “paying attention to”. Accordingly, being aware of – say – the color frame does not merely mean that one has the ability to distinguish between colors; rather, it means that one consciously distinguishes between colors *when thinking about the game*. (Just like you can hear someone talk and not listen to them, decision-makers filter out much of the data available to them as a way to reduce cognitive load.) To put it differently, an event of which an agent is unaware “is not necessarily one the agent *could not* conceive of, just one he *doesn't* think of at the time he makes his choice” (Dekel, Lipman, and Rustichini, 1998b, p. 524, italics in original). Thus, an event of which one is unaware is not the same as an event one has (thought about and) assigned probability zero: to see why the notions of unawareness and probability-zero are inherently different from each other, note that only an agent who had been *unaware* of an event might be affected by a question asking if that event may or may not happen.⁸

In summary, this paper contributes to the study of interactive unawareness by formally analyzing the role of *attribute awareness* in coordination games. We then put the theory to the test by experimentally investigating how choice behavior is impacted by changes in attribute awareness; we find that the best explanation of the data is consistent with our model and solution concept. The remainder is organized as follows: section II lays out the theory; section III introduces the experimental design; section IV contrasts the model’s predictions against a set of null hypotheses based on the standard Bayesian paradigm; section V presents the data, and section VI concludes.

II. Model and solution concept

1. *Actions and attributes*

Our model proposes that players identify an action with a (possibly partial) description of its observable attributes; we will see that this ultimately affects strategic deliberations. Although some of our insights can easily be generalized for a broad variety of problems with a coordination

⁸ For example, imagine that as you prepare to go out for a stroll, your partner asks you whether it might rain. If you had already assigned probability zero to the event “rain”, then you will be unmoved by such a question. If instead you had ignored that event – and you now notice it is cloudy outside – then you may react by taking an umbrella along.

element, for simplicity the analysis will center around a class of matching games where the action set consists of visual objects characterized by the “color”, “shape”, and “order” attributes. In this subsection, we shall define such a class of games and elaborate on the notion of attributes; in the next subsection, we will lay out the core components of the theory, namely, the players’ (partial) awareness of attributes and a novel solution concept.

Let Nature draw *three objects* from a 6-element set, one by one at random, and then place them in a column according to the order of selection, starting from the top. Denote by A the (unordered) 3-element collection of objects – henceforth “triplet” – selected by Nature. Next, define a game such that the *set of available actions* corresponds to A , and has generic member a ; assume that A becomes common knowledge as soon as Nature selects a triplet. For each player i, j and for each action $a \in A$, specify a player’s *payoff* to be a positive number π if players choose the same action a , and zero otherwise.

Each of the actions $a \in A$ is characterized by a tuple of attribute labels, as follows. Denote by K the *set of attributes*, with $K := \{C, S, O\}$, where C, S, O respectively denote the color, shape, and Nature’s order of selection/display of the objects. For each action (i.e., object) $a \in A$, assume that a ’s color, shape, and order attributes may each take on one value (“label”) – respectively – from $C = \{\text{cyan}, \text{lavender}, \text{turquoise}\}$, $S = \{\text{triangle}, \text{diamond}, \text{pentagon}\}$, and $O = \{\text{top}, \text{other}\}$. Intuitively, labels in C and S respectively represent the color and shape an object $a \in A$ may be, whereas labels in O indicate whether a is or is not the first object to be selected by Nature and placed at the top of the column.⁹ Finally, a *labeling function* $\ell: A \rightarrow \mathcal{L}$ formally specifies the relation between an action and its description, with \mathcal{L} being defined as $\mathcal{L} := C \times S \times O$; in a nutshell, a labeling function assigns to each action $a \in A$ a tuple of attribute labels $\ell(a) \in \mathcal{L}$.

EXAMPLE. The definition of a labeling function says that each $a \in A$ is characterized by a list of labels (one label per attribute). For instance, $\ell(a) = (\text{turquoise}, \text{diamond}, \text{top})$ completely identifies the object (i.e., action) whose color, shape, and order of selection/display are respectively turquoise, diamond, and top. Further, if the subsequent two objects being drawn by Nature consisted of a cyan diamond and a turquoise pentagon, then the set of available actions A would be described as $\ell(A) = \{(\text{turquoise}, \text{diamond}, \text{top}), (\text{cyan}, \text{diamond}, \text{other}), (\text{turquoise}, \text{pentagon}, \text{other})\}$.

⁹ Depending on the level of descriptive detail or the emphasis one puts on a particular position, the modeler may assume alternative specifications of the order attribute, such as $O'' = \{\text{middle}, \text{other}\}$ or $O''' = \{\text{bottom}, \text{other}\}$ or else $O'''' = \{\text{first}, \text{second}, \text{third}\}$. To keep the exposition simple, for the time being we assume that the order attribute is defined as $O \equiv O' = \{\text{top}, \text{other}\}$: as such, it limits itself to characterizing each available action $a \in A$ as either “top” or “not top”. Still, our experimental predictions and econometric analysis will account for all possible specifications of the order attribute.

2. Attribute awareness and frame-dependent rationalizability

So far, we have assumed that each of the actions $a \in A$ is characterized by a 3-tuple of attribute labels, under the implicit assumption that one is aware of (i.e., perceives and thinks about) all three attributes in K , with $K := \{C, S, O\}$. Below we will do away with this implicit assumption: in doing so, we will posit that a player may identify an action with a “partial description”, which involves solely those attributes (i.e., a subset of K) she currently perceives, while ignoring the rest.

a. PARTIAL DESCRIPTIONS. Before modeling the players’ awareness and beliefs, we must first introduce some notation to formalize how an action $a \in A$ is labeled when using just a subset of the attributes $K' \subseteq K$: to that purpose, we define the following system of projections.

For any subset $K' \in 2^K$ (where 2^K is the set of all subsets of K , i.e., the power set of K), let $\mathcal{L}_{K'}$ denote a collection of tuples of labels involving solely the attributes contained in K' . Below we refer to any subset $K' \in 2^K$ as a “frame”. For each such $K' \in 2^K$, a *partial description* $l_{K'}: A \rightarrow \mathcal{L}_{K'}$ is given by the commuting diagram:

$$\begin{array}{ccc} A & \xrightarrow{l} & \mathcal{L} \\ \text{identity}_A \downarrow & & \downarrow \text{projection} \\ A & \xrightarrow{l_{K'}} & \mathcal{L}_{K'} \end{array}$$

EXAMPLE. As in our earlier example, let $\ell(a) = (\textit{turquoise}, \textit{diamond}, \textit{top})$ identify the action $a \in A$ whose color, shape, and order of display are respectively turquoise, diamond, and top. Now, the commuting diagram above implies that if – say – $K' = \{C, S\}$, then it follows that $l_{\{C, S\}}(a) = (\textit{turquoise}, \textit{diamond})$: we say that this is a less expressive description than $\ell(a) = (\textit{turquoise}, \textit{diamond}, \textit{top})$, in that $l_{\{C, S\}}(a)$ involves solely a 2-tuple of (color, shape) attribute labels. As another example, note that if $K' = \{C\}$, then $l_{\{C\}}(a) = (\textit{turquoise})$: this is yet a less expressive description, as it involves just one attribute label. Finally, note that if $K' = \{\emptyset\}$, then $l_{\{\emptyset\}}(a) = (\text{blank})$: here the interpretation is that a is characterized trivially as a “nondescript object”.

Having defined a system of partial descriptions, we move on to articulate what it entails for a player to be aware of some such descriptions. To that end, we build on the epistemic structure proposed by Heifetz, Meier, and Schipper (2013a), which rests on a lattice of state spaces ordered by their “expressive scope” (i.e., ordered according to the extent to which the spaces account for any relevant contingencies).¹⁰ In particular, the remainder of the subsection is organized into the

¹⁰ Unlike the Heifetz et al. (2013a, 2013b) approach, which models awareness of actions per se, here we model awareness of the actions’ *attributes*. To do so, we shall define a state space for each subset of attributes (i.e., frame) $K' \in 2^K$. We will see that this construction allows us to restrict players’ beliefs in a natural way, which leads to a novel solution concept.

following parts: in (b) we embed partial descriptions into alternative *state spaces*; in (c) we model *awareness types*, which specify the frames that players perceive and think about; in (d) we define a novel solution concept, which we call *frame-dependent rationalizability*.




Figure 1 - A lattice structure defined by the power set of K , with an arrow connecting any two elements that are ordered via a superset relation \supseteq .

b. STATE SPACES. We set out to define a state space for each and every subset of attributes $K' \in 2^K$. Before doing so, we note that the members of 2^K can be (partially) ordered according to a superset relation \supseteq , thereby obtaining the lattice structure depicted in Figure 1. Each such $K' \in 2^K$ may be thought of as a different (as a whole) vocabulary with which to express facts; with that in mind, we now construct a system of multiple state spaces as follows.

For each $K' \in 2^K$ fix a *state space* $\Omega_{K'}$ – with generic member ω – in such a way as to obtain a *lattice of disjoint state spaces* $\{\Omega_{K'}\}_{K' \in 2^K}$, with the partial order given by the superset relation on the underlying attributes. The general interpretation is that state spaces higher up in the lattice involve a more expressive vocabulary, and thus provide a more thorough account of the possible states of the world. (Later on, we will see that each state space is interpreted as the particular viewpoint of a different player type, as defined in part (c) below.) Some qualifications are in order.

For any space $\Omega_{K'}$, each of the states $\omega \in \Omega_{K'}$ is assumed to encompass *all the relevant facts that can be expressed in terms of the K' frame* (i.e., expressible via the attributes contained in K'). Thus, each $\omega \in \Omega_{K'}$ implicitly includes player i 's potential K' -specific description of A , along with a description of opponent j 's frame K'' (with $K'' \subseteq K'$) and of j 's beliefs about i 's frame.¹¹ Whereas the paper's results apply to any lattice of state spaces satisfying this characterization, Figure 2 presents a specific lattice upon which our running example will be based.

¹¹ The construction implies that the upmost space (i.e., $\Omega_{\{C,S,O\}}$) provides the most exhaustive account of the resolution of all possible uncertainties in the model. Lower spaces have fewer states, as their expressive scope does not allow describing some contingencies; e.g., $\Omega_{\{C,S\}}$ does not account for the eventuality that some opponent frames may comprise the order attribute.




Figure 2 - A lattice of state spaces. Each of the rectangular boxes represents a K' -specific state space $\Omega_{K'}$: each such box is interpreted as the state space perceived by a different player type, as defined in part (c) below. Within any one box, each of the black dots represents a state ω : a state provides an account of all the relevant facts expressible in terms of the K' frame, such as i 's potential K' -specific description of A , plus a description of opponent j 's frame/beliefs. Each arrow points *from* a state *to* the corresponding information set perceived by a player type whose frame comprises one attribute less: such arrows are drawn for the sole benefit of the reader. (While player i knows how her state space maps to a less expressive space, i is unconcerned about any such “translation”, since her state space incorporates in itself *all* the relevant facts that i 's type can express about the game.)

NOTE: the upmost space $\Omega_{\{C,S,O\}}$ contains states where i 's frame K' comprises *color, shape, and order*. For visual clarity, this figure considers only 2 ordered triplets of objects, as indicated inside the upmost rectangular box; namely, $\{(turquoise, \diamond, 1), (cyan, \diamond, 0), (turquoise, \diamond, 0)\}$ and $\{(cyan, \triangleright, 1), (turquoise, \diamond, 0), (turquoise, \diamond, 0)\}$. A state obtains once Nature has drawn a triplet. States enclosed in the same oval account for the same description of A . Within any one oval, each state accounts for a different j type (i.e., different opponent frame/beliefs): such states are enclosed in the same oval to indicate that, once Nature has drawn a triplet, player i is uncertain as to the opponent's awareness. Spaces lower in the lattice have fewer states, since each $\Omega_{K'}$ accounts only for the eventuality of facing j types with frame K'' , for $K'' \subseteq K'$.

c. TYPE MAPPINGS. We now move on to develop the players’ unawareness belief structure by formalizing the notion of “awareness types”. In what follows, a player type is characterized by a mapping such that when ω obtains, the player may perceive and think about some other state(s).

Formally, denote the union of the above-defined state spaces by $\Omega := \cup_{K' \in 2^K} \Omega_{K'}$. For each i , then define a *type mapping* μ_i as a correspondence $\mu_i: \Omega \rightarrow \cup_{K' \in 2^K} \Delta(\Omega_{K'})$;¹² this is required to satisfy standard properties of unawareness structures (for discussion, see Heifetz, Meier, and Schipper, 2013a). In short, a type mapping μ_i assigns to each $\omega \in \Omega$ a set of probabilities about the states of which i is aware at ω . For instance, suppose that some state ω obtains: we interpret $\mu_i(\omega)(\{\omega'\})$ and $\mu_i(\omega)(\{\omega''\})$ respectively as the subjective probability about ω' and ω'' that i holds “at ω ” (i.e., when the true state is actually ω). In other words, when ω obtains, player i thinks that ω' or ω'' obtains, with probability $\mu_i(\omega)(\{\omega'\})$ and $\mu_i(\omega)(\{\omega''\})$, respectively. Note that ω' and ω'' may belong to a different (less expressive) space than the true state ω .

In a nutshell, $\mu_i(\omega)$ specifies a belief about each of the states that player i regards as possible when ω obtains. Whenever $\mu_i(\omega) \in \Delta(\Omega_{K'})$ for some $K' \in 2^K$, we informally refer to i as a K' type (i.e., i ’s current belief is concentrated on $\Omega_{K'}$); for any $K' \subset K$ we say that i is “unaware” of $K \setminus K'$.¹³

EXAMPLE. Suppose that $\omega \in \Omega_{\{C,S,O\}}$, and that the triplet drawn by Nature is describable as $\ell(A) = \{(cyan, triangle, top), (turquoise, diamond, other), (turquoise, pentagon, other)\}$. Here the true state $\omega \in \Omega_{\{C,S,O\}}$ corresponds to one of the eight states (i.e., black dots) on the *right*-hand-side of the upmost rectangular box, in Figure 2: in all of those eight states, player i ’s description of the triplet is given by $\ell(A)$, as indicated directly above the top right oval. Note that each of those states accounts for the eventuality of facing an opponent j with different frame/beliefs (i.e., different type): those eight states are enclosed in the same oval to indicate that i is aware of the drawn triplet’s description in terms of $\{C, S, O\}$, yet is uncertain about j ’s awareness, under the assumption that there are eight possible j types. (More generally, Figure 2 assumes that, for any drawn triplet, there is one state for each opponent frame that is expressible via K' : i.e., there is one state per $K'' \in 2^{K'}$; note that our results do not depend on this particular assumption.¹⁴) In summary, this example models the case where i is aware of all the attributes but is uncertain about j ’s awareness: accordingly,

¹² $\Delta(\Omega_{K'})$ denotes the set of probability measures over the state space $\Omega_{K'}$.

¹³ The symbol \setminus denotes set difference.

¹⁴ Thus, if $K' = \{C, S, O\}$, there is one state for each of the eight members of the power set of $\{C, S, O\}$; i.e., one state per $K'' \in \{\{C, S, O\}, \{C, S\}, \{C, O\}, \{S, O\}, \{C\}, \{S\}, \{O\}, \emptyset\}$. The interpretation is that – in each state – there is a j type that uses, respectively, *color-shape-order* labels, or only *color-shape* labels, or only *color-order* labels, etc.; note that as per our definition of a state, each of those states incorporates j ’s respective beliefs about i ’s frame. (In this connection, we note that prior research proved the existence of a universal unawareness type space containing all belief hierarchies: see Heifetz et al.’s 2011 working paper and Heinsalu, 2014.)

once Nature has drawn the triplet above, i 's type mapping will assign positive conditional probability to the aforementioned eight states, and probability zero to any other states in $\Omega_{\{C,S,O\}}$.

Next, consider the case where i 's type is $K' = \{S, O\}$, i.e., i is unaware of the color attribute. In terms of the model, this means that i 's type mapping will assign positive conditional probability to states in $\Omega_{\{S,O\}}$; formally, $\mu_i(\omega) \in \Delta(\Omega_{\{S,O\}})$. That is, i will perceive solely the states in $\Omega_{\{S,O\}}$ instead of the states in $\Omega_{\{C,S,O\}}$. More specifically, in the case of the triplet above, the possible states that i perceives are depicted by the four rightmost black dots in the second row of the lattice, denoted by $\omega^*, \omega^{**}, \omega^{***}, \omega^{****}$: in all of those four states, i 's description of the triplet is given by $\ell_{K'}(A) = \{(triangle, top), (diamond, other), (pentagon, other)\}$. To see why this information set has just four states, recall that – by construction – each $\Omega_{K'}$ includes only facts expressible via the attributes in K' : this implies that states in $\Omega_{K'}$ can account only for the eventuality of facing an opponent j with frame K'' , for $K'' \subseteq K'$. So, if i 's type is $K' = \{S, O\}$, then i 's perceived states will account only for opponent frames $K'' \in 2^{K'}$, that is, $K'' \in \{\{S, O\}, \{S\}, \{O\}, \emptyset\}$. The interpretation is that if i is unaware of the color attribute, i will solely think of states where j types describe A via, respectively, *shape-order* labels, or only *shape* labels, or only *order* labels, or *no* descriptive labels. Thus, if i is unaware of the color attribute, i will ignore any j types whose frames involve colors.

d. FRAME-DEPENDENT RATIONALIZABILITY. We now present a new solution concept. To do so, we build on Dekel, Fudenberg, and Morris' (2007) notion of interim correlated rationalizability, which captures interactions where there is a correlation between the state of the world and players' conjectures about the actions of others. Like the rationalizability notion in complete-information games (Bernheim, 1984; Pearce, 1984), Dekel et al.'s concept is defined via an iterated-deletion procedure. At each iteration, an action survives for a type only if: (i) it is a best-response to a belief assigning positive probability to type-action pairs of the opponents that have not yet been deleted; (ii) it is consistent with that type's beliefs about others and chance. Here, to allow for unawareness, we provide a notion of rationalizability whereby i does not best-respond to all the j types, but only to those types of which i is aware (i.e., as specified by $\mu_i(\omega)$). Another key novelty of our proposed solution concept is that it incorporates a pair of frame-dependent restrictions on beliefs, as follows.

The first restriction (“principle of indifference”) reflects Jacob Bernoulli's principle of insufficient reason, in the following sense: *if* one of the j types (frames) of which i is aware attaches the same label to two or more actions, then i believes that that j type will play those actions with the same probability. The second restriction (“oddity is prominence”) says that: *if* one of the j types (frames) of which i is aware attaches a distinct label to one action, and a common label to the other actions, then i believes that that j type will play the oddity with a higher probability.

Formally, for each player i , there is a (pure) strategy $s_i: \Omega \rightarrow A$ with the property that $\mu_i(\omega') = \mu_i(\omega) \Rightarrow s_i(\omega') = s_i(\omega)$, where ω, ω' denote two states that may or may not belong to the same space (i.e., this property implies that i 's strategy is a function of the set of states i perceives). Next, denote by $B_i(\omega)$ the *set of joint probability distributions* – with generic member β_i – over actions and states i perceives (when ω obtains). We interpret $\beta_i(\{a', \omega'\})$ as a belief i holds about the strategy j might take at ω' .

Given this, let $\beta_i(\{a', \omega'\}) \cdot u_i(a, a', \omega')$ denote player i 's expected utility from strategies $s_i(\omega') = a$ and $s_j(\omega') = a'$, where a, a' are generic members of A . Finally, define frame-dependent rationalizability (“FDR”) inductively as a sequence of *iterations* (indexed q) of beliefs B_i^q and strategies R_i^q of player i , as follows.

DEFINITION. For $\omega \in \Omega$ with $\mu_i(\omega) \in \Delta(\Omega_{K'})$, let $R_i^{q=0}(\omega) = A$. For $q > 0$,

$$B_i^q(\omega) := \left\{ \beta_i \in \Delta(A \times \Omega_{K'}) : \begin{array}{l} (0) \text{ } \textit{marg}_{\Omega_{K'}} \beta_i = \mu_i(\omega) \\ (1) \text{ } \mathbf{Principle of indifference:} \text{ if } \mu_j(\omega') \in \Delta(\Omega_{K''}) \text{ and } l_{K''}(a) = l_{K''}(a'), \\ \text{then } \beta_i(\{a, \omega'\}) = \beta_i(\{a', \omega'\}), \text{ for } a, a' \in A \\ (2) \text{ } \mathbf{Oddity is Prominence:} \text{ if } \mu_j(\omega') \in \Delta(\Omega_{K''}) \\ \text{and there exists an action } a'' \text{ such that} \\ l_{K''}(a'') \neq l_{K''}(a), l_{K''}(a'') \neq l_{K''}(a'), \text{ with } l_{K''}(a) = l_{K''}(a'), \\ \text{then } \beta_i(\{a'', \omega'\}) > \beta_i(\{a, \omega'\}) \text{ and } \beta_i(\{a'', \omega'\}) > \beta_i(\{a', \omega'\}), \text{ for } a, a', a'' \in A \\ (3) \text{ } \mathbf{Belief in (q - 1)-Rationality:} \beta_i(\{a, \omega'\}) > 0 \Rightarrow a \in R_j^{q-1}(\omega'), \text{ for } a \in A \end{array} \right\}$$

$$R_i^q(\omega) := \left\{ a \in R_i^{q-1}(\omega) : \begin{array}{l} \text{There exists } \beta_i \in B_i^q(\omega) \text{ such that} \\ a \in \arg \max_{a' \in A} \sum_{(a'', \omega') \in A \times \Omega_{K'}} \beta_i(\{a'', \omega'\}) \cdot u_i(a', a'', \omega') \end{array} \right\}$$

The set of player i 's frame-dependent rationalizable (“FDR”) strategies is

$$R_i(\omega) = \bigcap_{q=0}^{\infty} R_i^q(\omega).$$

Prior to illustrating the definition through a pair of results, a few comments are in order. Condition (0) above requires that i 's various beliefs be internally consistent, in the sense that the marginal of β_i over $\Omega_{K'}$ must be equal to $\mu_i(\omega)$. Conditions (1)-(2) define the above-discussed belief restrictions.^{15,16} Condition (3) is a standard feature of rationalizability concepts, requiring

¹⁵ As for (1), note that different formulations of Bernoulli's principle of “insufficient reason” have found application in previous models, like the Harsanyi and Selten (1988, pp. 70-74) principle of *invariance with respect to isomorphisms*, Crawford and Haller's (1990) *attainable strategies*, Bacharach's (1993) definition of *option sets* and related definitions in Casajus (2000), Janssen (2001), and Blume and Gneezy (2010). Unlike us, all these models use the principle of insufficient reason within the context of an equilibrium solution.

that (at each iteration q) i assign positive probability only to type-action pairs of the opponent that have not been deleted yet (i.e., that are rationalizable at $q - 1$). Finally, the set of i 's FDR strategies is defined as the collection of actions that survive an infinite sequence of iterations.

OBSERVATION 1. Consider a matching game where i 's type perceives a single attribute (besides \emptyset): assume that this attribute induces an oddity. If i 's type assigns any positive probability to a state where j is aware of the oddity, then the unique FDR strategy is for i to play that oddity.

Proof. The observation is easily proved by example. To that purpose, fix an arbitrary matching game and a lattice of state spaces. Without loss of generality, consider a triplet describable as $\ell(A) = \{(turquoise, diamond, top), (cyan, diamond, other), (turquoise, pentagon, other)\}$. Of the several possible states associated with this triplet [which in Figure 2 are depicted by the eight black dots on the *left-hand-side* of the upmost rectangular box], suppose that some $\omega \in \Omega_{\{C,S,O\}}$ obtains.

Consider an arbitrary K' comprising a single attribute, and assume that player i is a K' type: for instance, $K' = \{C\}$. Accordingly, when $\omega \in \Omega_{\{C,S,O\}}$ obtains, i 's type mapping will assign positive conditional probability solely to some states in $\Omega_{\{C\}}$; formally, $\mu_i(\omega) \in \Delta(\Omega_{\{C\}})$. [In Figure 2, the states of which i is aware are depicted by the two leftmost black dots in the third row of the lattice, denoted by $\hat{\omega}$ and $\tilde{\omega}$.] Note that each of those perceived states describes the action set as $\ell_{K'}(A) = \{(cyan), (turquoise), (turquoise)\}$: each such state accounts for one of the two opponent frames K'' that i can think of, with $K'' \in 2^{K'}$ (i.e., if i 's type is $K' = \{C\}$, then i can only think of $K'' \in \{\{C\}, \emptyset\}$).

We move on to apply our definition of FDR strategies to the game above. We begin with condition (0): this requires that i 's various beliefs be internally consistent, in the following sense. Consider one of i 's perceived states, say, $\hat{\omega}$. Denote by $\beta_i(\{a, \hat{\omega}\})$ i 's belief that the j type at $\hat{\omega}$ might play a . For any action set $A = \{a, a', a''\}$, condition (0) says that $\beta_i(\{a, \hat{\omega}\}) + \beta_i(\{a', \hat{\omega}\}) + \beta_i(\{a'', \hat{\omega}\})$ must be equal to $\mu_i(\omega)(\{\hat{\omega}\})$ (recall that the latter expression denotes the probability i assigns to $\hat{\omega}$).

In addressing conditions (1)-(2), let $\hat{\omega}$ refer to the state where j 's type is $K'' = \{C\}$. That is, according to player i , $\mu_j(\hat{\omega}) \in \Delta(\Omega_{\{C\}})$: this means that i believes that the j type at $\hat{\omega}$ describes A as $\ell_{K''}(A) = \{(cyan), (turquoise), (turquoise)\}$. Hence, condition (1) says that player i believes that the j type at $\hat{\omega}$ will play the two same-label (*turquoise*) actions with equal probability. (2) says that player i believes that this j type will play *cyan* with probability greater than each *turquoise* action.

Next, let $\tilde{\omega}$ refer to the state where j 's type is aware of no attributes, that is, $K'' = \{\emptyset\}$. Here, according to player i , $\mu_j(\tilde{\omega}) \in \Delta(\Omega_{\{\emptyset\}})$, in which case player i believes that the j type at $\tilde{\omega}$ describes A as $\ell_{K''}(A) = \{(nondescript\ object), ((nondescript\ object)), ((nondescript\ object))\}$. Hence, condition (1) says that player i believes that the j type at $\tilde{\omega}$ will play (all) same-label actions with equal probability. We further note that condition (2) has no bite here.

¹⁶ As for (2), note that the present formulation emphasizes simplicity over generality. Yet the idea behind the principle (i.e., a distinct minority stands out) is easily extended to games with more than three actions (e.g., the formulation may be generalized so that beliefs are inversely related to the frequency of each label). Incidentally, it is worth noting that Battigalli and Siniscalchi's (2003) notion of Δ -rationalizability incorporates belief restrictions that are similar in spirit to (but more general than) our approach; their concept however does not allow for unawareness, nor does it address prominence.

Based on the analysis above, at the first iteration (i.e., $q = 1$) the belief operator $B_i^1(\omega)$ includes any probability distribution β_i such that: (i) the j type at $\hat{\omega}$ is more likely to play *cyan*, and (ii) the j type at $\tilde{\omega}$ plays each of the three actions with equal probability. (Note: trivially any $B_i^1(\omega)$ satisfies condition (3), since by definition $R_j^0(\omega) = A$.) What about $R_i^1(\omega)$? It is easy to see that the only best-response to any such $\beta_i \in B_i^1(\omega)$ is for i to play *cyan*. This in turn implies that, owing to condition (3), $B_i^2(\omega)$ and successive iterations of the belief operator will remove from consideration any β_i such that the j type at $\hat{\omega}$ plays *cyan* with probability less than 1. To conclude, as long as player i regards $\hat{\omega}$ as possible (i.e., for any positive $\mu_i(\omega)(\{\hat{\omega}\})$), i 's expected utility is maximized by playing the color oddity. Thus, the only FDR strategy is for i to play that oddity. ■

Note that although the above behavioral prediction (that a player will “choose an oddity”) could be accounted for by other solution concepts, the underlying epistemic assumptions would differ greatly. In fact, our solution concept is characterized simply by common belief in rationality *and* in the two belief restrictions (principle of indifference, oddity is prominence). By contrast, solutions in the form of equilibrium refinements remain controversial as a predictive device, due to their demanding assumptions as to what the players must know to achieve coordination; also, note that normative models presume that the players' perception is fixed, and so they do not address the case where perception changes during the game. Conversely, our solution concept intuitively justifies how an increase in one's attribute awareness may lead to one's change in strategy.

OBSERVATION 2. Consider the same scenario as in Observation 1, but now suppose that before taking action i notices that [at least] one previously-ignored attribute induces another oddity. If the state where j is aware of a new oddity is assigned (by i 's updated type) a higher probability than the other states, then the unique FDR strategy is for i to play that new oddity.

Proof. Without loss of generality, suppose (as in Observation 1) that A is describable as $\ell(A) = \{(turquoise, diamond, top), (cyan, diamond, other), (turquoise, pentagon, other)\}$, and that some state $\omega \in \Omega_{\{C,S,O\}}$ obtains. Next, assume that before taking action, player i 's type K' changes from $\{C\}$ to, say, $\{C, S\}$. This increase in awareness implies that when $\omega \in \Omega_{\{C,S,O\}}$ obtains, player i 's (updated) type mapping will assign positive conditional probability to some states in $\Omega_{\{C,S\}}$; formally, $\mu_i(\omega) \in \Delta(\Omega_{\{C,S\}})$.¹⁷ [In Figure 2, the new possible states of which player i has become aware are depicted by the four leftmost black dots in the second row of the lattice, denoted by $\omega', \omega'', \omega''', \omega''''$.] Note: each of those perceived states accounts for one of the four opponent frames K'' of which i can think, with $K'' \in 2^{K'}$ (i.e., if i 's type is $K' = \{C, S\}$, then i can only think of $K'' \in \{\{C, S\}, \{C\}, \{S\}, \emptyset\}$).

¹⁷ The updated type mapping should reflect the prior mapping, when possible. For instance, the sum of the probabilities assigned to any states involving color (i.e., $K'' = \{C, S\}$ or $K'' = \{C\}$) *after* the change in awareness should be equal to the probability assigned to the one state involving color (i.e., $K'' = \{C\}$) *before* the change.

Per conditions (1)-(2), the belief operator $B_i^1(\omega)$ now includes any probability distribution β_i such that: (i) the j type (at ω'') with $K'' = \{C\}$ is more likely to play the color oddity, (ii) the j type (at ω''') with $K'' = \{S\}$ is more likely to play the shape oddity, and finally (iii) the j type (at ω'''') with $K'' = \emptyset$ is just as likely to play each of the three actions.¹⁸ In this case, an analogous line of reasoning as in Observation 1 implies that: if the state with $K'' = \{S\}$ is assigned a higher probability than the other states, i 's expected utility is maximized by playing the shape oddity (i.e., *pentagon*). If so, the only FDR strategy is for i to play this “new” oddity. ■

III. Experimental design

Below we put our model to the test. To do so, we present the following lab experiment.

At the beginning of the experiment, each person is assigned to a computer (the experiment was conducted using zTree; Fischbacher, 2007), and paired with an unknown partner. Participants are then told that everyone is being shown the same six objects on her own screen: specifically, participants see six “blocks” (i.e., colored geometric shapes), as described in Table 1 below; note that no such numbers or labels are shown to the subjects.

Object no.	1	2	3	4	5	6
Color	<i>Cyan</i>	<i>Cyan</i>	<i>Lavender</i>	<i>Lavender</i>	<i>Turquoise</i>	<i>Turquoise</i>
Shape	<i>Triangle</i>	<i>Diamond</i>	<i>Triangle</i>	<i>Pentagon</i>	<i>Diamond</i>	<i>Pentagon</i>

Table 1 - The six objects in the experimental game (object numbers/labels *not* shown to the subjects).

Initially the objects are loosely arranged in a hexagonal fashion (one per vertex), and collectively occupy the left-hand side of the screen. After each subject has viewed the six objects on her screen, the computer program selects three objects – one by one – by sliding them and placing them in a column (on the right-hand-side of the screen) according to the order of selection, starting from the top. (The three-object selection is identical for each participant; note that the experimental game, including its actions and attributes, reflects the formal model in section II.1.¹⁹)

¹⁸ Note that the beliefs about the j type with $K'' = \{C, S\}$ are *not* restricted by conditions (1)-(2). This is because a description in terms of multiple attributes involves three distinct labels (one color-shape label per action); that is, $\ell_{K''}(A) = \{(turquoise, diamond), (cyan, diamond), (turquoise, pentagon)\}$. Since no two labels are identical, conditions (1)-(2) have no bite.

¹⁹ To ensure that the set-up above be common information, a summary description of the game – including the following message – is read aloud by the experimenter: “The computer program will select three of the objects, and will then display those three objects to every participant in the same fashion and order.” Relatedly, we note that the design we presented in early working-paper versions of this study involved slightly different experimental instructions; following a reviewer’s suggestion, we edited the instructions to remove a possible ambiguity and re-ran the experiment, obtaining new (but qualitatively similar) data.

The rest of the objects subsequently disappear from the screen and subjects complete a few tasks, described below. In the end, each subject is asked to indicate her choice of one of the three objects, with the goal of coordinating with her partner in the pair. Each member of a pair receives a payoff of \$1.25 if each chooses the same object; each receives nothing otherwise.

Our between-subjects design includes the Baseline and All-Aware treatments. The list below details the entire sequence of events in the *Baseline* treatment.

- i. Each subject is assigned to a computer terminal and is shown the paper instructions.
- ii. Each subject is paired with an unknown partner. Subjects are presented six objects on their screen, three of which are later selected by the program and put in a column.
- iii. An on-screen message prompts subjects to label those three objects [“PART A”].
- iv. An on-screen message prompts subjects to estimate the probabilities of the three objects being chosen by others [“PART B”]; they are informed that good guesses will be rewarded with an additional payment.²⁰
- v. An on-screen message prompts subjects to choose an object by ticking the relevant box; they are reminded that their payoff will be \$1.25 if both members of the pair choose the same object, \$0 otherwise [“PART C”].
- vi. Steps *ii.-v.* are repeated for 9 more rounds, whereby in each round a new three-object selection is implemented by the computer program and shown to each member of a pair. (In each round, participants are randomly assigned to another pair and are so informed.) No feedback is given between rounds.
- vii. Payment.

A few comments are in order. We begin by noting that each of the ten rounds presents a different three-object selection, and so each round differs from the others in the characteristics of the available action set. (No feedback is given between rounds, that is, games.) To ensure that oddities be balanced across positions, our design implements pre-randomized blocks; this procedure also ensures that participants across sessions be presented with the same sequence of triplets.^{21,22}

²⁰ Subjects are shown a pie chart with three spokes, which they may adjust so that a sector’s relative area represents the estimated likelihood of an object being chosen by others (as an identifier, within each sector there is the label the subject entered at step *iii.*). We incentivized this task by informing subjects that if at least one of their three estimates differed by no more than 5 percentage points from the realized value, they would receive an extra payment of \$0.25 at the end of the session. No feedback was given before the end of the session. We note that this elicitation mechanism has a simplicity advantage over alternative mechanisms: as such, it minimizes the need for lengthy instructions. Like other mechanisms, this one has drawbacks too; e.g., subjects are directly incentivized to state (at least) one true estimate, which in some cases may cause slight deviations from truthful reporting.

²¹ At the beginning of every round, subjects are shown the same six blocks as in round 1, but three different blocks are subsequently selected in each round. The initial position of the six blocks is re-shuffled in every successive round, but it is identical for all subjects taking part in the same round.

²² More explicitly, the constrained pre-randomization ensures that: (*i*) color/shape oddities are *balanced across positions* (i.e., top, middle, and bottom of the column); (*ii*) the sequence of games is *representative of the theoretical distribution*

We stress that the objects are not pre-labeled, so that we do not impose exogenous frames. Still, to help subjects identify options when navigating across tasks, at the beginning of each round subjects are invited to type a short text (3-15 characters) in each of three boxes beside the objects. Note that the labeling is for the subject's reference only, and our hypotheses do not rely on such idiosyncratic strings (many of the strings are not clear-cut, and were we to use such free-form inputs in the analysis we would have to exercise our discretion in assigning a value to each alphanumeric string, undermining the tests' objectivity). In brief, our analysis revolves around the *guess* and *choice* data, respectively elicited in Part B and C of each round. For the instructions, see the Online Appendix.

The design of the *All-Aware* treatment is the same as the Baseline except for PART B, which presents three extra questions, as shown in the following transcript:

«Recall that – in Part C of the experiment – you will be prompted to pick one object in order to coordinate with your partner. Prior to that, we would like to know what you think about other participants in this room. Please answer the following questions by moving the sliders to the desired percentages. Note: your partner will not be asked to answer these questions.

- 1) How likely do you think it is that the other participants have noticed the order in which the objects have been drawn by the computer program? Please move the below slider...
- 2) How likely do you think it is that the other participants have noticed the different colors of the objects? Please move the below slider...
- 3) How likely do you think it is that the other participants have noticed the different shapes of the objects? Please move the below slider...»

Note that the order in which questions 1-3 are presented is randomized in each round. Subjects enter their beliefs by moving a slider (i.e., one slider for each question) to the desired percentage, with the slider ranging from 0% to 100%. Note that the purpose of the questions is to make subjects *privately* aware of multiple attributes: to that end, All-Aware participants are matched into pairs with Baseline participants; accordingly, we inform All-Aware participants that their (Baseline) counterparts are *not* exposed to questions 1-3 above. This feature of the design ensures that All-Aware participants believe – correctly – that their partners' awareness has not been raised exogenously. (Obviously, to keep the Baseline participants' awareness unchanged, they are not informed about the extra questions to which their All-Aware counterparts are exposed.)

of triplets, as follows. Given the attributes in Table 1, the theoretical probability that Nature randomly selects a triplet with *no* color or shape oddities is 1/10 (see objects {1,4,5} and {2,3,6} in Table 1). Furthermore, the probability of a triplet containing one *color* oddity is 1/3 (e.g., {1,2,6}); the probability of a triplet containing one *shape* oddity is 1/3 (e.g., {2,3,5}); the probability of a triplet containing *both* color and shape oddities is 1/3 (e.g., {4,5,6}).

After presenting the questions above, the All-Aware treatment proceeds to the task described at step *iv.* of the Baseline. The rest of the treatment is identical to the Baseline.

Prior to stating our hypotheses, we note that it is easy for one to distinguish objects according to an attribute *if* one thinks about that attribute. That is, (barring some rare eye disorder) one could fail to see differences in the objects only if one did not pay attention to the attributes: such an unconscious neglect corresponds to the game-theoretic notion of unawareness. Per the epistemic literature, “being aware of an event” means that the *event is taken into account when making a decision* (Modica and Rustichini, 1999, p. 274). Hence, being aware of – say – the color frame does not mean that one can generally distinguish between colors; rather, it means that one consciously distinguishes between colors when thinking about the game. So, an event of which an agent is unaware “is not necessarily one the agent *could not* conceive of, just one he *doesn't* think of at the time he makes his choice” (Dekel et al., 1998b, p. 524, italics in original).

IV. Experimental hypotheses

We now show that our model produces numerous intuitive predictions regarding the treatments above; such predictions will be articulated in the form of alternative hypotheses, while null hypotheses will be based on the “standard” single-state-space Bayesian paradigm (i.e., incomplete-information models without unawareness). In doing so, we spell out how our model’s assumptions are at odds with the default economic model of knowledge.²³

We start by considering predictions that relate specifically to the All-Aware treatment. Recall that the All-Aware treatment manipulation involves a sequence of pre-play questions about whether others have or have not noticed the three (color, shape, and order) attributes of the currently drawn triplet. Those questions may be viewed as tautologies, as in “ E is the case or $\neg E$ is the case”, where E and $\neg E$ respectively represent the event such that j *has* and *has not* noticed an attribute.

Now, in a Bayesian game one is assumed to always know the full set of states; given the standard models’ underlying properties, then for any event E (i.e., for any subset of states) one is always aware of E and $\neg E$ (Dekel et al., 1998a). This implies that one’s awareness level would be unaffected by the questions above. Further, note that Bayesian games presume common knowledge of the information structure: this precludes i from believing that j is uncertain about something,

²³ Dekel et al. (1998a, Theorems 1, 2, pp. 166-169) prove that standard models of knowledge preclude unawareness (the case where one does not know an event *and* does not know that one does not know it): the rough intuition is that if an agent knows the full set of states, then she cannot be unaware of any events; as we shall see, the null hypotheses below follow from this result. In analyzing the data (in section V), we will also address early models with “non-standard” information structures, such as Bacharach (1993).

without j knowing so. In other words, a player cannot conceive of the opponent unconsciously ignoring something. Accordingly, our first null hypothesis (H1) is that All-Aware participants do not think that their Baseline counterparts may overlook features of the game such as the actions' attributes.²⁴

H1. All-Aware participants believe that their Baseline counterparts notice with certainty any differences in the objects with respect to colors, shapes, and order.

By contrast, our alternative hypothesis is that “*All-Aware participants believe that others may fail to notice some features of the action set, with such beliefs varying across players as well as games*”. To see how this alternative hypothesis stems from our model, note that we circumvent the Dekel et al. impossibility result (that standard models preclude unawareness) by building on Heifetz et al.'s (2006, 2013a) system of multiple state spaces; thus, here the full information structure is not common knowledge. With regard to the experiment, our model says that each All-Aware participant personifies a player with $\mu_i(\omega) \in \Delta(\Omega_{\{C,S,O\}})$ who (perceives all attributes and) is uncertain about the opponent's perception. In short, our model's alternative hypothesis is that All-Aware participants consider it possible that their counterparts might not notice every feature of the actions; so, their beliefs about whether others perceive some attribute (henceforth “*awareness beliefs*”) may each be less than 100%. (Instead, barring the rare case where people may be affected by an untreated eye disorder, standard models imply that any such elicited beliefs should be close to 100%.²⁵)

We proceed to H2. This is a three-fold hypothesis (with components *a*, *b*, *c*) concerning the relation between awareness beliefs and behavior. Following up on the discussion above, note that standard models imply no particular relation between one's strategy and one's awareness beliefs (since each elicited belief ≈ 1); so, they entail the next null hypotheses.

²⁴ Bayesian models for games with incomplete information (Harsanyi, 1967–68) formally capture *conscious* uncertainty. That is, a situation where a player knows that she cannot distinguish elements of the space of uncertainty. Such a space is assumed to be commonly known, and may include *opponents' strategies* or *moves by Nature* (or both); note that – in the context of our experiment – the latter case would correspond to a player who knows that she cannot tell apart the triplets, hence the actions, as she cannot see well. (Such a model would aptly represent a situation where it is common knowledge that players have an untreated eye disorder; however, it is an implausible characterization of our experiment, due to the low occurrence of any such impairments. In fact, the common type of color blindness consists of a decreased ability to tell green from red: this cannot affect our games, since all the objects involve shades of blue.)

²⁵ This and all other *null hypotheses* below are based on a standard model, defined as follows: consider a coordination game as described in the experimental instructions; next, assume that – as the game begins – players become commonly aware that the actions vary in color, shape, and order; finally, fix an (arbitrary) epistemic type space, whereby each type is associated with a hierarchy of beliefs about the players' behavior conditional on the drawn triplet. Note that this is a Bayesian game (without unawareness) where the space of uncertainty consists of the opponents' actions, given the observed triplet.

H2.a. The frequency with which All-Aware participants choose *color* oddities is not related to their belief about others noticing the different colors (for brevity, this belief is hereafter referred to as μ_i^C).

H2.b. The frequency with which All-Aware participants choose *shape* oddities is not related to their belief about others noticing the different shapes (henceforth μ_i^S , for brevity).

H2.c. The frequency with which All-Aware participants choose *top, middle, or bottom* objects is not related to their belief about others noticing the order of the objects (henceforth μ_i^O , for brevity).

In brief, null hypothesis H2.a above applies to All-Aware participants in games containing color oddities. The alternative hypothesis to H2.a is that “*the frequency of play of color oddities is positively related to μ_i^C , and not to μ_i^S or μ_i^O* ”. To see how this alternative hypothesis stems from our model, recall that player i 's FDR strategy corresponds to the oddity induced by the attribute most-likely noticed by the opponent (see Observation 2 for a proved statement). In particular, when i believes that others are more likely to notice the different colors (relative to shapes and order, i.e., when $\mu_i^C \geq \mu_i^S$ and $\mu_i^C \geq \mu_i^O$), i should choose a color oddity if there is one. Therefore, to test whether a color oddity is indeed i 's FDR strategy, the experimenter must verify that the frequency of play of color oddities is positively related to μ_i^C , and not to μ_i^S or μ_i^O .

We turn to H2.b, which applies to All-Aware participants in games with shape oddities. Thus, following the same reasoning as above, our model's alternative hypothesis to H2.b is that “*the frequency of play of shape oddities is positively related to μ_i^S , and not to μ_i^C or μ_i^O* ”.

We move on to address H2.c, which concerns the All-Aware participants' propensity to choose the n -th option. Prior to elaborating on our test, we note that whereas color and shape attributes induce a natural labeling, the same is not necessarily true of the order attribute. In fact, depending on the level of descriptive detail or the emphasis one puts on a particular position, the order attribute could be associated with any of the following labelings: *i.* ((*top*), (*other*), (*other*)); *ii.* ((*other*), (*middle*), (*other*)); *iii.* ((*other*), (*other*), (*bottom*)); *iv.* ((*first*), (*second*), (*third*)).²⁶ While labeling *iv* does not generate “order oddities”, *i*, *ii*, *iii* respectively pull a subject toward the top, middle, and bottom blocks (thus possibly away from any color or shape oddities). To simplify the exposition, we have so far identified the order attribute with the first labeling (as per footnote 9); yet in conducting the data analysis we shall account for all the labelings above. Accordingly, our

²⁶ Note that each such labeling has analogous translations that entail the same partition. For instance, labeling *iv* is equivalent to ((*top*), (*middle*), (*bottom*)), and to ((*rishon*), (*sheni*), (*shlishi*)), and to ((*primo*), (*secondo*), (*terzo*))...

model’s alternative hypothesis to H2.c is that “*the frequency of play of (just) one of the three positions is positively related to μ_i^O , whereas no position is positively related to μ_i^C or μ_i^S* ”.²⁷

In the following, we contrast the All-Aware treatment with the Baseline. Recall that the treatments are identical except for the three extra questions asked of All-Aware participants, which may be viewed as tautologies (i.e., “ E is the case or $\neg E$ is the case”). For any event E , remember that in a standard model one is always aware of E and $\neg E$: it follows that one cannot learn (from the All-Aware questions) anything that one did not know already. Standard models therefore predict that the All-Aware questions have no impact on participants’ behavior (regardless of the solution concept). So, there should be no behavioral differences across treatments, as in the following null hypothesis.²⁸

H3. Average choices do not vary between the Baseline and All-Aware treatments.

By contrast, our model’s alternative hypothesis is that “*average choices vary across treatments*” due to an increase in attribute awareness, resulting from the three extra questions that we asked of All-Aware participants (i.e., if one had not been aware of E and $\neg E$ in the first place, then the question itself would automatically generate awareness of those events). To see how this would impact choices, we note that the randomized assignment of subjects to either treatment guarantees a-priori similar samples across treatments (thus, *a-priori similar* choice distributions, on average). Now recall that – from the experimenter’s perspective – the All-Aware questions ensure that each participant in the All-Aware treatment will shift (from an unobservable type) to type $K' = \{C, S, O\}$. If a participant’s prior type was different, in that before seeing those questions she had ignored some attribute/s, then she might reconsider her strategy as a result of the updated type. This leads to behavioral differences across treatments. (See Claim 1 in the Appendix for a formal statement.)

The next hypothesis addresses whether any between-treatment differences in game play are reflected in the subjects’ *guesses* about which objects will be chosen by others. Such guesses were elicited from both All-Aware and Baseline participants (see task *iv.*, in section III), and should not

²⁷ Recall that our design involves a pre-randomization mechanism, ensuring that color and shape oddities be balanced across positions. So, high beliefs μ_i^O cannot be driven by an abundance of color/shape oddities in a particular position. As we will see, the data reject null hypothesis H2.c in favor of our model’s alternative hypothesis: more precisely, the frequency of play of the top object is positively related to μ_i^O , which points to a labeling such as $((top), (other), (other))$.

²⁸ Also note that the All-Aware treatment manipulation hints at multiple attributes at once (without directing subjects’ attention to any one frame in particular), thereby minimizing any implicit demand effects on the part of our questions.

be confused with the above-discussed awareness beliefs. Due to the same reasoning as above, standard models entail the following null hypothesis.

H4. Average guesses (about which objects will be chosen by others) do not vary across the Baseline and All-Aware treatments.

Conversely, our model’s alternative hypothesis is that “*average guesses vary across treatments*” due to an increase in attribute awareness on the part of All-Aware participants.

Our final hypothesis has to do with differences in *coordination rates* between treatments. Once again, if participants were unaffected by the All-Aware questions (as is implied by standard Bayesian models), then average choices would not vary across treatments; if so, coordination rates would not vary either. This means that standard models entail the following null hypothesis.

H5. Coordination rates do not vary between the <Baseline, Baseline> and <Baseline, All-Aware> pairs of subjects.

By contrast, our model’s alternative hypothesis is that “*the treatment manipulation causes a decrease in coordination rates*”, as follows. As a benchmark, take the hypothetical case in which Baseline participants are paired with Baseline participants (i.e., <Baseline, Baseline> pairs); we then compare such a reference group with the case in which Baseline participants are paired with All-Aware participants, as in our experimental design (i.e., <Baseline, All-Aware> pairs).²⁹ Here, the model implies that coordination rates for <Baseline, All-Aware> pairs must be weakly lower than those for <Baseline, Baseline> pairs. The informal argument is that some Baseline participants overlook an attribute, so they will not match their All-Aware counterparts who choose an object based on that attribute; generalizing, increases in attribute awareness affect the variance of the choice distributions, thus coordination rates. (See Claim 2 in the Appendix for a formal statement.)

V. Experimental results

1. *General procedures, recruitment, and earnings*

Experimental sessions were conducted at UCSB, with subjects being recruited from a broad range of academic departments via ORSEE (Greiner, 2015). A total of 108 subjects took part in our

²⁹ Note that it would not make sense to consider the case in which All-Aware participants are paired with each other, since that would contradict the information provided during the experiment (i.e., recall that All-Aware participants believe – correctly – that their counterparts’ awareness has not been raised exogenously).

six sessions. Subjects on average earned a total payoff of about \$12 (over ten games, including a \$5 show-up fee), with minimum (maximum) earnings of \$6.50 (\$17.50). On average, sessions had 18 subjects and lasted about 40 minutes. In each session, *half* of the participants was assigned to the All-Aware treatment and *half* to the Baseline. No subject could participate in more than one session.

2. Tests of H1

We begin by analyzing the All-Aware treatment. Here we address the null hypothesis H1, concerning the distribution of subjective probabilities about whether others *do* or *do not* notice an attribute (collectively referred to as “awareness beliefs”, for short). To that end, Figure 3 reports histograms for the awareness beliefs with respect to *colors*, *shapes*, and *order* (which for brevity we respectively denote by $\mu_i^C, \mu_i^S, \mu_i^O$, where each variable is a number belonging to the interval [0,1]).³⁰




Figure 3 - All-Aware treatment. The first, second and third panel (from the top) show histograms for the beliefs about others noticing differences in the objects with respect to colors, shapes and order, respectively. Note: the data refer to *per-subject mean beliefs* elicited (as percentages) over ten games; such beliefs are treated as continuous variables, with the width of each histogram bin covering about 5 percentage points.

³⁰ Here, such elicited beliefs fully determine a type mapping, as defined in section II.2. For instance, let’s denote by $\tilde{\omega}$ the state where j is aware of the color attribute *only* (i.e., j is aware of the color attribute, and *not* of the shape or order attributes). Then, the probability i ’s type mapping assigns to $\tilde{\omega}$ is computed as $\mu_i^C \cdot (1 - \mu_i^S) \cdot (1 - \mu_i^O)$.

A quick glance at Figure 3 shows that All-Aware participants are really *not certain* that their Baseline counterparts would notice every feature of the action set, as is instead implied by the Bayesian paradigm (i.e., incomplete-information models presume common knowledge of the information structure, which precludes a player from believing that the counterpart may overlook any features of the game). In brief, we see from Figure 3 that only about 30% of the color (1st panel) and shape (2nd panel) awareness beliefs are within the 95-100 percent interval, and less than 15% of the order (3rd panel) awareness beliefs are within that interval. Notably, the rate of *order* awareness beliefs that fall within the 95-100 percent interval is significantly different ($N = 108$ obs., $z = 2.053$, $p = 0.040$, two-tailed test of proportions) than the corresponding rate for *color* and *shape* awareness beliefs, confirming that these responses are not merely noise.

Now, to reject the null hypothesis H1 (“All-Aware participants believe that their Baseline counterparts notice with certainty any differences with respect to colors, shapes, and order”), we just need to show that the awareness beliefs are less than 100 percent. So, with 70%-85% of the distributions *outside* of the 95-100 percent interval, we can readily reject H1. People do not have faith that everything will be observed, giving some scope for unawareness to have an impact. Even a conservative Wilcoxon signed-rank test for whether the median belief differs from the value of 95 percent (i.e., instead of 100 percent, thus allowing for “almost certainty”) is strongly significant: for *color* awareness beliefs ($N = 54$ obs., $z = -4.376$, $p = 0.000$, two-tailed), for *shape* awareness beliefs ($N = 54$ obs., $z = -4.280$, $p = 0.000$, two-tailed), and finally for *order* awareness beliefs ($N = 54$ obs., $z = -5.997$, $p = 0.000$, two-tailed). Therefore, All-Aware participants are far from certain that their Baseline counterparts would notice every feature of the action set. (Incidentally, we stress that the tests above are conducted on the sample of per-subject mean beliefs, to satisfy the assumption of independence of observations; i.e., the tests use one observation per participant.) We conclude that the data reject the null hypothesis H1.

Moving on, we note that evidence against H1 is consistent with our model as well as with earlier theories featuring “non-standard” information structures, such as Bacharach’s (1993) *variable frame theory* and related work (Bacharach and Stahl, 2000; Casajus, 2000; Janssen, 2001), henceforth collectively referred to as VFT. While each VFT variant differs somewhat from others, they each allow for some heterogenous awareness. (As we previously noted, our model draws on this literature, yet unlike VFT we define a lattice of state spaces and provide a new solution concept, which avoids the drawbacks resulting from some VFT assumptions.) In particular, VFT assumes

that – in a given population – the probability that someone thinks of an attribute is a constant, and it is independent of the “physical distribution of the objects; and of features of the player, such as her skill, her experience of similar games, and her opportunities to search for descriptors” (Bacharach and Bernasconi, 1997, p. 13). In the context of our experiment, this implies that the probability of noticing, say, color differences is a constant, and it is invariant to the specific color of the objects across games. Further, VFT posits that players’ beliefs consistently reflect that probability.³¹

Since the VFT assumptions above are too strong to generate testable predictions, here we relax them. Instead of assuming that any individuals who perceive an attribute k somehow hold identical beliefs, we assume that beliefs are concentrated around a value that we take to be the mean of a normal distribution. We first report summary statistics to have a sense of whether *mean beliefs are similar across players*; we then test if those beliefs come from the same normal distribution. Descriptive statistics relating to the (per-subject) mean dataset confirm that *color* awareness beliefs are fairly dispersed across players: $avg. = 86.825$, $s.d. = 12.552$; further, the Shapiro-Wilk W test (i.e., a common test for normality) rejects the hypothesis that subjects’ beliefs come from the same normal distribution ($N = 54$ obs., $z = 4.723$, $p = 0.000$). For *shape* awareness beliefs, we find that $avg. = 83.470$, $s.d. = 15.912$ and, again, the Shapiro-Wilk W test rejects the hypothesis that beliefs come from the same normal distribution ($N = 54$ obs., $z = 3.779$, $p = 0.000$). For *order* awareness beliefs, we find that $avg. = 71.140$, $s.d. = 18.587$; here, the test provides mild evidence against the hypothesis that beliefs come from the same normal distribution ($N = 54$ obs., $z = 1.413$, $p = 0.078$), however the standard deviation of the observed distribution is even higher than in previous cases. So, while VFT implies that – for each attribute k – beliefs should not vary across k -perceiving individuals, our data show that they vary substantially and do not exhibit normality.

The tests so far utilized per-subject mean observations. In what follows, instead, we test if *an individual’s own beliefs are similar across the sequence of games* (recall that, per VFT, the probability that someone thinks of an attribute is a constant, and it is independent of the physical distribution of the objects). To that end, we report a Friedman test (i.e., the non-parametric analog to

³¹ This assumption is motivated by the notion of the “acquisition of mutual beliefs among normal agents”. For example, suppose that in a given population the actual probability that players normally notice color differences is p . There, VFT posits that *all* (Bacharach, 1993, pp. 260-263) or *some* (Bacharach and Stahl, 2000, p. 230) of the color-perceiving players correctly know p , while the rest of the players do not hold any beliefs about color perception. In the context of our experiment, this implies that All-Aware participants generally hold correct beliefs about their Baseline counterparts. Yet, this is implausible in games where one has no experience about the others’ perceptual limitations. In designing an empirical test of VFT, the late Bacharach acknowledged this and other issues, adding that “[the authors] wanted to test VFT constructively, that is, in such a way that evidence of incorrectness would show how to improve it” (Bacharach and Bernasconi, 1997, p. 12).

the Repeated Measures ANOVA) conducted on the entire sequence of awareness beliefs, consisting of ten games per subject. This test indicates significant differences across games in the case of *color* awareness beliefs ($\chi_9^2 = 22.946, p = 0.006$), *shape* awareness beliefs ($\chi_9^2 = 19.087, p = 0.024$), and *order* awareness beliefs ($\chi_9^2 = 37.465, p = 0.000$). So, the data do not support the VFT prediction that beliefs reflect the fact that the probability of noticing an attribute must be constant across games.

In summary, the data reject null hypothesis H1, thereby contradicting the predictions of a standard model; moreover, we find that the data contradict some VFT predictions. In fact, we find support for our model’s alternative hypothesis, that is: “*All-Aware participants believe that others may fail to notice some features of the action set, with such beliefs varying across players as well as games*”.

3. Tests of H2

We proceed to test null hypothesis H2 (with components *a, b, c*), which as usual is based on a standard model (per footnote 25). In brief, the single-state-space Bayesian paradigm implies no particular relation between one’s strategy and one’s awareness beliefs: in fact, in that model one is aware of everything, and knows that others are aware of everything (thus, each elicited belief ≈ 1).

	[I] Choice of the <i>color</i> oddtity	[II] Choice of the <i>shape</i> oddtity	[III] Choice of the <i>top</i> object	[IV] Choice of the <i>middle</i> object	[V] Choice of the <i>bottom</i> object
Belief about others noticing <i>colors</i> μ_i^C	.029 *** (.010)	-.020 ** (.008)	-.014 (.013)	-.001 (.015)	.024 (.015)
Belief about others noticing <i>shapes</i> μ_i^S	-.012 ** (.006)	.021 ** (.008)	.001 (.008)	.002 (.010)	-.004 (.009)
Belief about others noticing <i>order</i> μ_i^O	-.013 ** (.006)	.004 (.006)	.008 * (.004)	-.001 (.005)	-.010 *** (.003)
<i>Round</i>	-.034 (.036)	.261 ** (.129)	.042 (.052)	-.040 (.085)	-.013 (.114)
<i>Constant</i>	-.317 (.966)	-1.88 (1.158)	.289 (1.034)	-.818 (.892)	-2.164 ** (1.045)
Pseudo R ²	0.038	0.077	0.017	0.003	0.024
Obs.	324	324	540	540	540

Table 2 - Logit model coefficients, with robust standard errors adjusted for two-way clustering on the subjects and the games. *, ** and *** respectively indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$ for the relevant Z-statistic, two-tailed tests. Note: models [I] and [II] use All-Aware data from any games with at least a *color* and a *shape* oddity, respectively; model [III]-[V] use All-Aware data from all the ten games. For details about the distribution of oddities across games, see footnote 22.

Specifically, H2.a says that there is no relationship between color awareness beliefs μ_i^C and choice behavior in *games containing color oddities*. For a formal test, column [I] of Table 2 presents a logit model with a subject's choice of the color oddity as the binary dependent variable; the list of predictors includes each of the awareness beliefs μ_i^C , μ_i^S , μ_i^O , as well as a round variable (controlling for any time effects). Note that model [I] uses All-Aware data from any games containing a color oddity, with half such games containing a shape oddity as well; in order to provide a most conservative test, we present robust standard errors adjusted for two-way clustering (Cameron, Gelbach, and Miller, 2011), on the subjects and on the games. Results from model [I] show a significant positive effect of μ_i^C (i.e., the belief about others noticing the different colors) on the likelihood of playing the color oddity; also, we find a significant negative effect of μ_i^S and μ_i^O . Hence, the data reject H2.a in favor of the alternative hypothesis derived from our model: “*the frequency of play of color oddities is positively related to μ_i^C , and not to μ_i^S or μ_i^O* ”.

We turn to null hypothesis H2.b, which posits no relationship between shape awareness beliefs μ_i^S and choice behavior in *games containing shape oddities*. In short, column [II] of Table 2 presents a logit model with a subject's choice of the shape oddity as the binary dependent variable; as before, predictors include the awareness beliefs as well as a round variable. Note that model [II] uses All-Aware data from any games containing a shape oddity, with half such games containing a color oddity as well; again, robust standard errors are adjusted for clustering. Model [II] indicates a significant positive effect of μ_i^S (i.e., the belief about others noticing the different shapes) on the likelihood of playing the shape oddity; additionally, we find a significant negative effect of μ_i^C and no significant effect of μ_i^O . So, the data reject H2.b in favor of our model's alternative hypothesis: “*the frequency of play of shape oddities is positively related to μ_i^S , and not to μ_i^C or μ_i^O* ”.

We move on to null hypothesis H2.c, which concerns the *full sample of games* (as opposed to a subsample of games containing color or shape oddities). This hypothesis posits no relationship between order awareness beliefs μ_i^O and behavior. Instead, our model predicts that “*the frequency of play of (just) one of the three positions is positively related to μ_i^O , while no position is positively related to μ_i^C or μ_i^S* ”. (In this regard, it is worth noting that our design involves a pre-randomization mechanism, ensuring that color and shape oddities be balanced across positions: this implies that high beliefs μ_i^O cannot be driven by an abundance of color/shape oddities in a particular position.) That said, columns [III]-[V] of Table 2 present logit models respectively consisting of a subject's choice of the top, middle, and bottom object as the binary dependent variable; again, predictors include the awareness beliefs, along with a round variable (the models use All-Aware data from each of the ten games, with robust standard errors adjusted

for clustering). Model [III] shows a mildly significant positive effect of μ_i^O (i.e., the belief about others noticing the order of the objects) on the likelihood of playing the top object. Interestingly, while none of the models [III]-[V] indicates any significant effect of μ_i^C or μ_i^S , model [V] shows a significant negative effect of μ_i^O on the likelihood of playing the bottom object. We conclude that the data reject H2.c in favor of our model’s alternative hypothesis: specifically, we find evidence that *the frequency of play of the top object is positively related to μ_i^O* , which points to a labeling such as $((top), (other), (other))$.

To recap, the tests reject H2 (a, b, c). In fact, the tests support our rationalizability concept, according to which i ’s FDR strategy is the oddity induced by the attribute most-likely noticed by j .

4. Tests of H3

In the rest of the paper, we contrast the All-Aware treatment with the Baseline. We start by testing null hypothesis H3, which pertains to the distributions of choice data: based on the standard Bayesian paradigm, this null hypothesis says that there should be no behavioral differences across the two treatments, on average (i.e., if one never overlooks any attributes, then the three All-Aware questions cannot alter one’s view of the game). Instead, our model’s alternative hypothesis is that choices vary between treatments due to an increase in attribute awareness on the part of All-Aware participants.

Prior to discussing our tests, we shall present some summary statistics. To that end, we let $((a_1), (a_2), (a_3))$ denote a generic, ordered triplet of objects (a_n refers to the n -th available option). In the Baseline, subjects chose a_1 , a_2 and a_3 , respectively, 40.18%, 33.15% and 26.67% of the time (averaging across ten games); in the All-Aware treatment, instead, 50.00%, 26.30% and 23.70%, respectively. These average (across-games) distributions hint at differences between treatments, yet these statistics are uninformative as to the relation between one’s attribute awareness and the objects’ attributes *in each game*. So, like before, we must run an analysis of the full sample of individual observations (adjusting standard errors for clustering on the subjects and the games).

Columns [I]-[III] of Table 3 present logit models respectively consisting of a subject’s choice of a_1 , a_2 and a_3 as the dependent variable. The list of predictors includes: (i) “*Treat*”, a treatment indicator taking on value 0 or 1 when a subject is assigned to the Baseline or All-Aware, respectively; (ii) “*Odd*”, a dummy for whether the object a_n specified in the column’s header is a color/shape oddity or not; (iii) a *Treat*Odd* interaction variable; (iv) “*Else*”, a dummy for whether the triplet of objects contains another color/shape oddity (i.e., other than the object a_n specified in the column’s header); (v) a “*Round*” variable, controlling for any time effects.

	[I] Choice of a_1	[II] Choice of a_2	[III] Choice of a_3
<i>Treat</i> (Treatment indicator)	.575 ** (.282)	-.623 ** (.279)	-.341 (.239)
<i>Odd</i> (Dummy for whether the object a_n specified in the column's header is a color/shape oddity or not)	1.315 *** (.360)	.986 * (.579)	1.343 *** (.503)
<i>Treat*Odd</i> (Interaction variable)	-.429 (.299)	.371 (.362)	.284 (.317)
<i>Else</i> (Dummy for whether the triplet contains another color/shape oddity or not)	-1.215 *** (.177)	-.904 ** (.364)	-.243 (.426)
<i>Round</i>	-.026 ** (.011)	.001 (.080)	.018 (.045)
<i>Constant</i>	.299 (.202)	-.758 (.809)	-1.591 *** (.548)
Pseudo R ²	0.095	0.117	0.096
Obs.	1080	1080	1080

Table 3 - Logit model coefficients, with robust standard errors adjusted for two-way clustering on the subjects and the games. *, ** and *** respectively indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$ for the relevant Z -statistic, two-tailed tests. Note: $((a_1), (a_2), (a_3))$ denotes a generic, ordered triplet of objects (a_n refers to the n -th available option); each model uses data from all the ten games. The indicator *Treat* takes on value 0 or 1 when a subject is respectively assigned to the Baseline or All-Aware.

Results from Table 3 show a positive effect of *Odd*, and a negative effect of *Else*. This means that the frequency of play of an object a_n increases when that object is a color/shape oddity, and it decreases when another object is a color/shape oddity. Remarkably, the significance of *Treat* in models [I]-[II] provides evidence that behavior varies across treatments, in such a way that All-Aware participants are more likely to choose a_1 (and less likely to choose a_2) than Baseline participants. We interpret such differences as being driven by an increase in the awareness of the order labeling $((top), (other), (other))$.

Further, we note that the *interaction* variable in each of the models of Table 3 is non-significant, indicating that All-Aware participants are as likely as Baseline participants to choose color and shape oddities (counted together). This appears to suggest that our treatment manipulation did not affect the All-Aware participants' consideration of color/shape frames. However, the non-significance of the interaction might conceal opposite-sign changes in the frequency of play of color and shape oddities. To check that, we analyze the subsample of games that contain both a color oddity and a shape oddity at once (in the same triplet of objects). Given this subsample, we consider a simple logit model with one's

choice of color oddities as the binary dependent variable, and the treatment indicator as the sole predictor: such a model reveals a significant negative impact of the All-Aware treatment on the choice of color oddities ($N = 324$, $\text{coef.} = -.275$, $z = -2.20$, $p = 0.028$, two-tailed logit with standard errors adjusted for clustering). Given the same subsample, we also consider a logit model with one's *choice of shape oddities* as the binary dependent variable, and the treatment indicator as the sole predictor: this model shows a positive non-significant effect of the All-Aware treatment on the choice of shape oddities ($N = 324$, $\text{coef.} = .268$, $z = 0.98$, $p = 0.325$, two-tailed logit with clustered standard errors).

The analysis above rejects H3, thereby supporting our model's alternative hypothesis that "*average choices vary across treatments*". In short, the data indicate behavioral differences pointing to an increase in the awareness of a labeling like (*top*), (*other*), (*other*). In other words, the All-Aware manipulation affects choices, because it makes participants think about order (and to a lesser extent shape) frames to which they would otherwise not have paid attention.

5. Tests of H4

We now verify if the above differences in game play are reflected in the subjects' *guesses* about which objects will be chosen by others. Such guesses were elicited from both All-Aware and Baseline participants (see task *iv.*, in section III), and should not be confused with the above-discussed awareness beliefs. As usual, the null hypothesis is based on the standard Bayesian paradigm, which entails no differences in the guesses stated across treatments, on average.³² On the other hand, our model's alternative hypothesis is that guesses vary across treatments due to an increase in attribute awareness.

Columns [I]-[III] of Table 4 present OLS models consisting of a subject's guess about respectively a_1 , a_2 and a_3 being chosen by others, as the dependent variable; the list of predictors is the same as in Table 3. Results from Table 4 confirm a significant positive effect of *Odd*, and a negative effect of *Else*: this means that guesses (about a_n being chosen by others) increase when the object is a color/shape oddity, and they decrease when another object is a color/shape oddity. Further, the mild significance of *Treat* in model [I] provides some evidence that guesses about a_1 (which corresponds to the *top* object) change – namely, increase – in the All-Aware treatment

³² Indeed, since participants are randomly assigned to either treatment, the large sample size and the homogeneity in the make-up of the subject pool give us *no* reason to assume *prior* differences across samples. Given this, if subjects' awareness is unaffected by the treatment manipulation (in that subjects pay attention to all the frames regardless of the treatment, as is assumed by Bayesian models), then Bayesian models imply no differences in the distribution of guesses.

relative to the Baseline. At the same time, the significant negative effect of *Treat* in model [III] confirms that guesses about a_3 (the *bottom* object) decrease in the All-Aware treatment relative to the Baseline. Finally, note that the *interaction* variable in model [I] reveals that the All-Aware participants’ guesses about a_1 increase relatively less when a_1 is a color/shape oddity.³³

	[I] Guess about a_1	[II] Guess about a_2	[III] Guess about a_3
<i>Treat</i> (Treatment indicator)	3.891 * (2.269)	-1.866 (1.270)	-2.938 ** (1.496)
<i>Odd</i> (Dummy for whether the object a_n specified in the column’s header is a color/shape oddity or not)	6.777 *** (1.578)	5.989 ** (2.507)	6.080 ** (2.564)
<i>Treat*Odd</i> (Interaction variable)	-2.669 *** (.603)	1.262 (2.235)	2.706 (2.520)
<i>Else</i> (Dummy for whether the triplet contains another color/shape oddity or not)	-6.488 *** (.371)	-4.991 *** (1.547)	-1.325 (1.375)
<i>Round</i>	.095 (.111)	-.221 (.290)	.141 (.207)
<i>Constant</i>	37.650 *** (.824)	34.663 *** (2.759)	29.254 *** (2.225)
R ²	0.068	0.115	0.088
Obs.	1080	1080	1080

Table 4 - OLS regression coefficients, with robust standard errors adjusted for two-way clustering on the subjects and the games. *, ** and *** respectively indicate $p < 0.10$, $p < 0.05$ and $p < 0.01$ for the relevant Z-statistic, two-tailed tests. Note: $((a_1), (a_2), (a_3))$ denotes a generic, ordered triplet of objects (a_n refers to the n -th available option); each model uses data from all the ten games. The indicator *Treat* takes on value 0 or 1 when a subject is respectively assigned to the Baseline or All-Aware.

We conclude that the results reject H4 and support our model’s alternative hypothesis that “*average guesses vary across treatments*”. In particular, the present analysis corroborates our earlier interpretation that All-Aware participants exhibit an increase in the awareness of the order labeling $((top), (other), (other))$. Still, it is worth discussing a speculative counter-argument, according to which the All-Aware manipulation could lead subjects who had already been aware of all three attributes to somehow “reassess” their guesses (i.e., despite there being no change in attribute awareness). In this

³³ If a_1 is a color/shape oddity, the predicted value of the guess about a_1 goes from 40.70% to 41.93% when moving from the Baseline to the All-Aware treatment. By contrast, if a_1 is *not* a color/shape oddity, the predicted value of the guess about a_1 increases from 32.58% to 36.47% when moving from Baseline to All-Aware. The latter change is consistent with an increase in the awareness of a labeling such as $((top), (other), (other))$, since color/shape oddities are *not* a factor in that case.

respect, we note that while such a reassessment might be justified at the margin (were one to attach nearly the same probability to multiple objects), one cannot justify a *significant* reassessment *unless* new information has arisen. Relatedly, we stress that the All-Aware treatment manipulation hints at several attributes at once (without directing subjects' attention to any one frame in particular), thereby minimizing any implicit demand effects. Thus, subjects could *not* interpret a question in the All-Aware treatment as a signal that "one should particularly pay attention to an attribute", since such a signal would be uninformative in that the three questions are symmetric to each other. This falsifies the argument that one may observe a significant change in the distribution of guesses even without a change in awareness.

6. Tests of H5

Here we test null hypothesis H5, which concerns *coordination rates* across the Baseline and All-Aware treatments. While Bayesian models imply no between-treatment differences in behavior and so no difference in coordination rates, the alternative hypothesis derived from our model is that the treatment manipulation causes a change, whereby coordination rates for <Baseline, All-Aware> pairs are lower than those for <Baseline, Baseline> pairs (see Appendix for a formal claim).

We start by reporting some summary statistics. In keeping with previous studies, we report expected coordination rates (as opposed to actual frequencies of coordination), computed at the session level.³⁴ As a benchmark, we consider the hypothetical case in which a Baseline participant is paired with another Baseline participant, which yields a 50.2% coordination rate, averaging across sessions. We then turn to the case in which an All-Aware participant is paired with a Baseline participant – as per our experimental design – which yields a 44.6% coordination rate (averaging across sessions). This provides informal evidence against the null hypothesis.

For a formal test, we now contrast the *distribution of per-session* mean coordination rates for <Baseline, All-Aware> pairs, against the distribution for <Baseline, Baseline> pairs. Such a comparison allows us to conclude that the coordination rates differ significantly from each other under a Wilcoxon signed rank sum test ($N = 12$ obs., $z = -2.201$, $p = 0.027$, two-tailed).

Additionally, a binomial test can help us verify if (same-direction) differences across coordination

³⁴ Actual coordination rates depend on individual choices and on a stochastic element, that is, the random assignment of partners: in a relatively small sample, this random element might bias the rates. "Thus, the *actual* frequency of coordination has no special significance; it is more appropriate to consider the *expected* frequency of coordination" (Mehta, Starmer, and Sugden, 1994, p. 663, italics in original). Here we compute the expected coordination rates from the observed distribution of individual choices, by calculating the probability that two subjects match in a certain round (i.e., triplet) and session; per-session mean rates are then computed by averaging across all the rounds in a session.

rates are due to chance, as opposed to treatment-induced variations in awareness. In this regard, we note that <Baseline, All-Aware> pairs exhibit lower coordination rates than <Baseline, Baseline> pairs in six out of six sessions, confirming that variations in subjects' awareness do have an impact on coordination ($N = 6, p = \left(\frac{1}{2}\right)^5 = 0.031$, two-tailed). In brief, the data reject null hypothesis H5 in favor of our model's alternative hypothesis that "*the treatment manipulation causes a decrease in coordination rates*". As per our prediction, an increased awareness can hurt coordination.

VI. Discussion and conclusion

We have proposed a model allowing for heterogeneity in players' awareness as to the attributes of the action set. We have then provided a test of competing hypotheses about the impact of frames on choice behavior. The data confirm that changes in attribute awareness cannot be plausibly accounted for by a standard Bayesian model (which precludes unawareness in the first place); in fact, we find that the best explanation of the data is consistent with our proposed model and solution concept.

Taken together, the present theory and evidence provide a coherent account of the impact of varying multi-attribute awareness. This account builds on research streams such as the formal analysis of unawareness (e.g., Dekel et al., 1998a; Heifetz et al., 2013a) and the study of labelings (Bacharach, 1993; Bacharach and Stahl, 2000; Janssen, 2001). With reference to the latter, we note that while Bacharach's work and related variants allowed for heterogeneous awareness – as previously discussed – their solution concepts rely on quite strong assumptions; thus, they are not well suited as a predictive device in experimental games where one has no experience about the others' perceptual limitations, or in cases where one's perception may change during the game. In reviewing some of these drawbacks, the late Bacharach hinted at several directions for future research, some of which we have taken up here (e.g., see the discussion in Bacharach and Bernasconi, 1997, pp. 12-13). Relatedly, we note that we depart from Bacharach and Bernasconi's seminal experiment, since their design did not attempt to raise subjects' awareness, nor did it test if subjects best-respond to their "awareness beliefs". (For an overview of related experiments, see Rojo Arjona, 2020.)

In this connection, Blume and Gneezy (2010) studied an interesting form of heterogeneity in frame awareness, drawing on the influential analysis of symmetries by Crawford and Haller (1990) and Blume (2000). Specifically, Blume and Gneezy's experimental design presupposes two levels of cognitive ability (low and high), where only the high type perceives a symmetry in the structure of the options that induces a Pareto-efficient strategy. While the Blume-Gneezy design did not elicit awareness beliefs, their results show evidence that high types conceive of and react to low types

(e.g., subjects behaved differently when trying to match an unknown partner rather than their own previous choice). Summing up, in the authors' words, the Blume-Gneezy design aims to address "logical inference" and "mathematical foci" (2010, p. 490); as such, it does not address individuals that are similar in their cognitive skills, and yet heterogenous in the extent to which they simply *think about* multiple symmetric attributes when making a decision.

In the domain of (normative) equilibrium theories, we note that Arad and Rubinstein's (2019) model does address multi-attribute thinking (in Colonel Blotto games); however, their model pertains to problems where "players have accumulated experience in playing the game and have settled on a particular mode of behavior" (p. 286). More explicitly, their work does not aim to address the case where inexperienced subjects suffer from partial perception: hence, unawareness has no role in their model. In a similar vein, Alós-Ferrer and Kuzmics (2013) study players with common knowledge of a focal frame; so, their model "is not a descriptive one of how players behave in a given framed game in the lab [...], but rather how players should and perhaps eventually will behave, after generations of teaching and learning" (p. 229).

In the experimental realm, research on the impact of "prescriptive frames" has compared behavior across problems with exogenously-assigned *game labels* (e.g., "Community game" vs. "Wall Street game"; Kay and Ross, 2003) or *action labels* (e.g., {cooperate, defect} vs. {out, in}; Andreoni, 1995).³⁵ There, the experimenter induces a prescriptive frame by evoking individualistic versus cooperative norms, which direct subjects' attention toward the action one *ought to* take in the context evoked. In this respect, we note that our design does not involve any exogenous labels; also, our design evokes multiple frames at once, thereby minimizing any implicit demand effects.

To conclude, this paper has studied how attribute awareness relates to rational choice; the data confirm that changes in attribute awareness do affect choice behavior in matching games. Going forward, we note that since our solution concept is defined inductively, it can also be used to generate predictions for finite, low levels of mutual belief in rationality (e.g., as in the Level-*k* literature) in games without a pure coordination motive. Indeed, a better appreciation of the perception-action link may lead to new applications for games incorporating an element of conflict (Shah and Ludwig, 2016). Everyday life does in fact show that our mental framing can influence which actions we consider – and then choose – in a broad array of interactions.

³⁵ For evidence on exogenous labels, see O'Neill (1987), Rapoport and Boebel (1992), Rubinstein, Tversky, and Heller (1997), Tversky (2004). For evidence on labels in the context of Level-*k* reasoning, see Crawford and Iriberri (2007), Crawford, Gneezy, and Rottenstreich (2008), and Sontuoso and Bhatia (2021); see also Alaoui, Janezic, and Penta (2020).

References

- Alaoui, Larbi, Katharina A. Janezic and Antonio Penta.** 2020. "Reasoning About Others' Reasoning" *Journal of Economic Theory*, 189: 105091.
- Alós-Ferrer, Carlos and Christoph Kuzmics.** 2013. "Hidden Symmetries and Focal Points" *Journal of Economic Theory*, 148(1): 226-258.
- Andreoni, James.** 1995. "Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments" *Quarterly Journal of Economics*, 110(1): 1-21.
- Arad, Ayala and Ariel Rubinstein.** 2019. "Multidimensional Reasoning in Games: Framework, Equilibrium, and Applications" *American Economic Journal: Microeconomics*, 11(3): 285-318.
- Arrow, Kenneth J.** 1974. *The Limits of Organization*. WW Norton & Company.
- Bacharach, Michael.** 1993. "Variable Universe Games" in *Frontiers of Game Theory*, ed. K. Binmore et al. Cambridge, MA: MIT Press.
- Bacharach, Michael and Michele Bernasconi.** 1997. "The Variable Frame Theory of Focal Points: An Experimental Study" *Games and Economic Behavior*, 19(1): 1-45.
- Bacharach, Michael and Dale O. Stahl.** 2000. "Variable-Frame Level-n Theory" *Games and Economic Behavior*, 32(2): 220-246.
- Battigalli, Pierpaolo and Marciano Siniscalchi.** 2003. "Rationalization and Incomplete Information" *B. E. Journal of Theoretical Economics: Advances*, 3(1): article 3.
- Bernheim, B. Douglas.** 1984. "Rationalizable Strategic Behavior" *Econometrica*, 52(4): 1007-1028.
- Blume, Andreas.** 2000. "Coordination and Learning with a Partial Language" *Journal of Economic Theory*, 95(1): 1-36.
- Blume, Andreas and Uri Gneezy.** 2010. "Cognitive Forward Induction and Coordination Without Common Knowledge: An Experimental Study" *Games and Economic Behavior*, 68(2): 488-511.
- Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller.** 2011. "Robust Inference with Multiway Clustering" *Journal of Business & Economic Statistics*, 29(2): 238-249.
- Casajus, André.** 2000. "Focal Points in Framed Strategic Forms" *Games and Economic Behavior*, 32(2): 263-291.
- Chen, Yi-Chun, Jeffrey C. Ely and Xiao Luo.** 2012. "Note on Unawareness: Negative Introspection Versus AU Introspection (and KU Introspection)" *International Journal of Game Theory*, 41(2): 325-329.
- Crawford, Vincent P., Uri Gneezy and Yuval Rottenstreich.** 2008. "The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures" *American Economic Review*, 98(4): 1443-1458.
- Crawford, Vincent P. and Hans Haller.** 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games" *Econometrica*, 58(3): 571-595.
- Crawford, Vincent P. and Nagore Iriberry.** 2007. "Fatal Attraction: Saliency, Naivete, and Sophistication in Experimental Hide-and-Seek Games" *American Economic Review*, 97(5): 1731-1750.
- Dekel, Eddie, Drew Fudenberg and Stephen Morris.** 2007. "Interim Correlated Rationalizability" *Theoretical Economics*, 2 (1): 15-40.
- Dekel, Eddie, Barton L. Lipman and Aldo Rustichini.** 1998a. "Standard State-Space Models Preclude Unawareness" *Econometrica*, 66(1): 159-173.
- Dekel, Eddie, Barton L. Lipman and Aldo Rustichini.** 1998b. "Recent Developments in Modeling Unforeseen Contingencies" *European Economic Review*, 42(3-5): 523-542.
- Fagin, Ronald and Joseph Y. Halpern.** 1987. "Belief, Awareness, and Limited Reasoning" *Artificial Intelligence*, 34(1): 39-76.
- Fischbacher, Urs.** 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments" *Experimental Economics*, 10(2): 171-178.
- Gauthier, David.** 1975. "Coordination" *Dialogue*, 14(2): 195-221.
- Greiner, Ben.** 2015. "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE" *Journal of the Economic Science Association*, 1(1): 114-125.
- Gunning, Robert C.** 2018. *An Introduction to Analysis*. Princeton University Press.
- Harsanyi, John C.** 1967-68. "Games with Incomplete Information Played by 'Bayesian' Players, Parts I-III" *Management Science*, 14(3): 159-182; 320-334; 486-502.
- Harsanyi, John C. and Reinhard Selten.** 1988. *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.

- Heath, Chip and Nancy Staudenmayer.** 2000. "Coordination Neglect: How Lay Theories of Organizing Complicate Coordination in Organizations" *Research in Organizational Behavior*, 22: 153-191.
- Heifetz, Aviad, Martin Meier and Burkhard C. Schipper.** 2006. "Interactive Unawareness" *Journal of Economic Theory*, 130(1): 78-94.
- Heifetz, Aviad, Martin Meier and Burkhard C. Schipper.** 2011. "Unawareness, Beliefs, and Speculative Trade" *Working Papers 47, University of California, Davis, Department of Economics*.
- Heifetz, Aviad, Martin Meier and Burkhard C. Schipper.** 2013a. "Unawareness, Beliefs, and Speculative Trade" *Games and Economic Behavior*, 77(1): 100-121.
- Heifetz, Aviad, Martin Meier and Burkhard C. Schipper.** 2013b. "Dynamic Unawareness and Rationalizable Behavior" *Games and Economic Behavior*, 81: 50-68.
- Heinsalu, Sander.** 2014. "Universal Type Structures with Unawareness" *Games and Economic Behavior*, 83: 255-266.
- Janssen, Maarten C. W.** 2001. "Rationalizing Focal Points" *Theory and Decision*, 50(2): 119-148.
- Kay, Aaron C. and Lee Ross.** 2003. "The Perceptual Push: The Interplay of Implicit Cues and Explicit Situational Construals on Behavioral Intentions in The Prisoner's Dilemma" *Journal of Experimental Social Psychology*, 39(6): 634-643.
- Kets, Willemien and Alvaro Sandroni.** 2021. "A Theory of Strategic Uncertainty and Cultural Diversity" *Review of Economic Studies*, 88(1): 287-333.
- Kreps, David M.** 1990. "Corporate Culture and Economic Theory" in *Perspectives on Positive Political Economy*, ed. J. E. Alt and K. A. Shepsle. Cambridge: Cambridge University Press.
- Mehta, Judith, Chris Starmer and Robert Sugden.** 1994. "The Nature of Salience: An Experimental Investigation of Pure Coordination Games" *American Economic Review*, 84(3): 658-673.
- Milgrom, Paul R. and John Donald Roberts.** 1992. *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall.
- Minsky, Marvin.** 1975. "A Framework for Representing Knowledge" in *The Psychology of Computer Vision*, ed. P. Winston. New York: McGraw-Hill.
- Modica, Salvatore and Aldo Rustichini.** 1994. "Awareness and Partitional Information Structures" *Theory and Decision*, 37(1): 107-124.
- Modica, Salvatore and Aldo Rustichini.** 1999. "Unawareness and Partitional Information Structures" *Games and Economic Behavior*, 27(2): 265-298.
- Okhuysen, Gerardo A. and Beth A. Bechky.** 2009. "Coordination in Organizations: An Integrative Perspective" *Academy of Management Annals*, 3(1): 463-502.
- O'Neill, Barry.** 1987. "Nonmetric Test of The Minimax Theory of Two-Person Zerosum Games" *Proceedings of the National Academy of Sciences*, 84(7): 2106-2109.
- Pearce, David G.** 1984. "Rationalizable Strategic Behavior and The Problem of Perfection" *Econometrica*, 52(4): 1029-1050.
- Rapoport, Amnon and Richard B. Boebel.** 1992. "Mixed Strategies in Strictly Competitive Games: A Further Test of The Minimax Hypothesis" *Games and Economic Behavior*, 4(2): 261-283.
- Rojo Arjona, David.** 2020. "Focal Points in Laboratory Experiments" in *New Palgrave Dictionary of Economics*, ed. M. Vernengo, E. Perez Caldentey and B. J. Rosser. London: Palgrave Macmillan.
- Rubinstein, Ariel, Amos Tversky and Dana Heller.** 1997. "Naive Strategies in Competitive Games" in *Understanding Strategic Interaction*, ed. W. Albers et al. Berlin: Springer-Verlag.
- Schank, Roger C. and Robert P. Abelson.** 1977. *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schelling, Thomas C.** 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Schipper, Burkhard C.** 2014. "Unawareness—A Gentle Introduction to Both The Literature and The Special Issue" *Mathematical Social Sciences*, 70: 1-9.
- Shah, Anuj K. and Jens Ludwig.** 2016. "Option Awareness: The Psychology of What We Consider" *American Economic Review: Papers & Proceedings*, 106(5): 425-29.
- Sontuoso, Alessandro and Sudeep Bhatia.** 2021. "A Notion of Prominence for Games with Natural-Language Labels" *Quantitative Economics*, 12(1): 283-312.
- Sugden, Robert.** 1995. "A Theory of Focal Points" *Economic Journal*, 105(430): 533-550.
- Tversky, Amos.** 2004. *Preference, Belief, and Similarity: Selected Writings*. Cambridge, MA: MIT Press.

APPENDIX

Claim 1. An increase in attribute awareness (on the part of some All-Aware participants) causes behavioral differences across treatments, *ceteris paribus*.

Proof. Assume that, in a given game, there is a subset of Baseline participants that does not think about all three attributes. For instance, and without loss of generality, suppose that some Baseline participants are aware of a single attribute (besides \emptyset); formally, there are some Baseline participants whose K' type is either $\{C\}$ or $\{S\}$ or $\{O\}$. It follows from Observation 1 that each such participant will choose the color, shape, and order oddity (if any), respectively. Now, regardless of the specific distribution of types, note that the randomized assignment of subjects to either treatment ensures *a-priori similar* samples (i.e., it ensures that the unobserved a-priori distribution of types be similar) across treatments. Then, it follows from Observation 2 that an increase in attribute awareness causes a change in strategy on the part of some of the All-Aware participants, relative to their Baseline counterparts. Such a change implies behavioral differences across treatments. ■

Claim 2. An increase in attribute awareness causes a fall in coordination rates, *ceteris paribus*.

Proof. We start by noting that distributions of choice data (on which coordination rates depend) can be represented as 3-dimensional vectors, where the n -th element of the vector indicates the frequency of play of the n -th object in a given sample of participants. Let the vector \mathbf{Y} denote the choice distribution of a sample of Baseline participants. Further, let $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Y}}^t$ respectively denote the “ex ante” and “ex post” choices of a sample of All-Aware participants; that is, $\tilde{\mathbf{Y}}$ represents the *unobserved* distribution of choices absent the treatment manipulation (i.e., prior to being exposed to the three All-Aware questions), whereas $\tilde{\mathbf{Y}}^t$ is the *observed* distribution of choices (i.e., after the treatment manipulation). Even though the experimenter does not observe $\tilde{\mathbf{Y}}$, note that the randomized assignment of subjects to either treatment ensures a-priori similar samples, across treatments. This implies that the vectors \mathbf{Y} and $\tilde{\mathbf{Y}}$ must be similar; formally, $\cos(\mathbf{Y}, \tilde{\mathbf{Y}}) \sim 1$, with \cos denoting the cosine similarity of the vectors.* Next, note that the expected coordination rate can be defined as the dot product of the vectors (i.e., the sum of the products of the vectors’ corresponding entries). Thus, in the hypothetical case in which untreated subjects are paired with each other, the coordination rate is given by $\mathbf{Y} \cdot \tilde{\mathbf{Y}}$, where $\mathbf{Y} \cdot \tilde{\mathbf{Y}} = \sum_{n=1}^3 Y_n \tilde{Y}_n$. Similarly, denote by $\mathbf{Y} \cdot \tilde{\mathbf{Y}}^t$ the coordination rate for the case in which *untreated* subjects are paired with *treated* subjects.

* The cosine similarity between (3-dimensional) vectors \mathbf{Y} and $\tilde{\mathbf{Y}}$ is defined as $\cos(\mathbf{Y}, \tilde{\mathbf{Y}}) = \frac{\sum_{n=1}^3 Y_n \tilde{Y}_n}{\|\mathbf{Y}\| \|\tilde{\mathbf{Y}}\|}$, where the numerator is the *dot product of the vectors* and the denominator is the *product of the Euclidean norm of each vector*. The reader can easily verify that, for any two vectors with non-negative values, such a similarity measure ranges between 0 and 1. In our case this has a nice interpretation. The cosine similarity is in fact 0 if every participant in a sample chooses object a and every participant in the other sample chooses object a' , for any $a' \neq a$. By contrast, the cosine similarity is 1 whenever the choice distributions are identical across the two samples.

That said, we want to compare $\mathbf{Y} \cdot \tilde{\mathbf{Y}}$ with $\mathbf{Y} \cdot \tilde{\mathbf{Y}}^t$ and determine which coordination rate is greater. To this purpose, we can use a well-known result from linear algebra and rewrite each of the dot products as follows. By the law of cosines (Gunning, 2018, p. 65), write $\mathbf{Y} \cdot \tilde{\mathbf{Y}} \equiv \|\mathbf{Y}\| \cdot \|\tilde{\mathbf{Y}}\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}})$; likewise, write $\mathbf{Y} \cdot \tilde{\mathbf{Y}}^t \equiv \|\mathbf{Y}\| \cdot \|\tilde{\mathbf{Y}}^t\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}}^t)$, with $\|\cdot\|$ denoting the Euclidean norm of a vector. Note that the coordination rates are now expressed as the product of three non-negative scalars, and hence we can divide each of the expressions by $\|\mathbf{Y}\|$. That leaves us to compare $\|\tilde{\mathbf{Y}}\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}})$ with $\|\tilde{\mathbf{Y}}^t\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}}^t)$. Before doing so, note that per Observation 2 an increase in attribute awareness causes a change in strategy on the part of some treated subjects. Under the assumption that participants' deviations from the ex-ante choice distribution are equally likely across objects, then the variance of $\tilde{\mathbf{Y}}^t$ will weakly decrease relative to $\tilde{\mathbf{Y}}$. (Note: the variance of a 3-object choice distribution is lowest, with a value of 0, when the frequency of play of each object is 1/3, whereas the variance is highest when every participant in the sample chooses the same one object.) It is easy to see that whenever $\text{var}(\tilde{\mathbf{Y}}) \geq \text{var}(\tilde{\mathbf{Y}}^t)$, then it must be that $\|\tilde{\mathbf{Y}}\| \geq \|\tilde{\mathbf{Y}}^t\|$. Since by assumption $\cos(\mathbf{Y}, \tilde{\mathbf{Y}}) \sim 1$, then it follows that $\|\tilde{\mathbf{Y}}\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}}) \geq \|\tilde{\mathbf{Y}}^t\| \cdot \cos(\mathbf{Y}, \tilde{\mathbf{Y}}^t)$. Hence, $\mathbf{Y} \cdot \tilde{\mathbf{Y}} \geq \mathbf{Y} \cdot \tilde{\mathbf{Y}}^t$. That is, coordination rates for <Baseline, Baseline> pairs must be (weakly) greater than those for <Baseline, All-Aware> pairs. ■

ONLINE APPENDIX

The Doors of Perception: Theory and Evidence of Frame-Dependent Rationalizability

Gary Charness *and* Alessandro Sontuoso

Experimental instructions (*Baseline* and *All-Aware* treatments)

NOTE: The screenshots below refer to both treatments, unless otherwise noted.

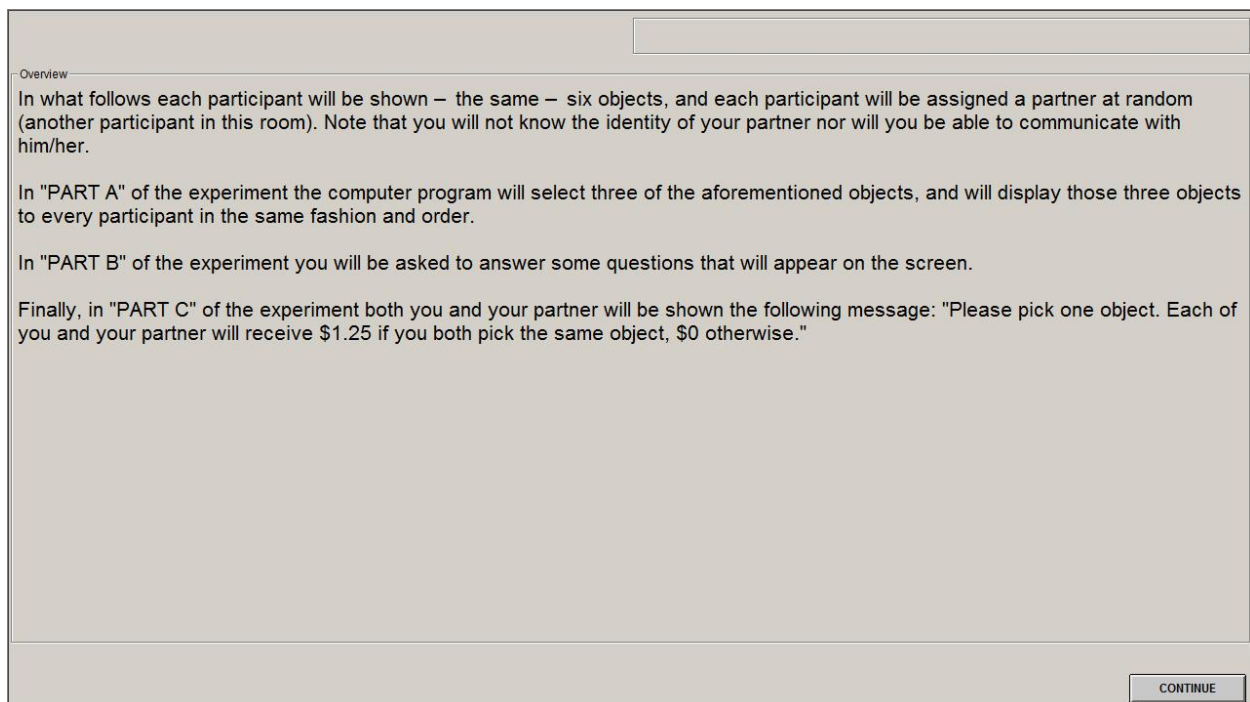
→Home screen:

Consent form.


→Screenshot 2 (transcript):

Hello and welcome to a decision-making experiment. You will receive a show-up fee, and can also earn additional money. The additional payment will be determined by your own choices and those made by some other participant, according to the rules described below.

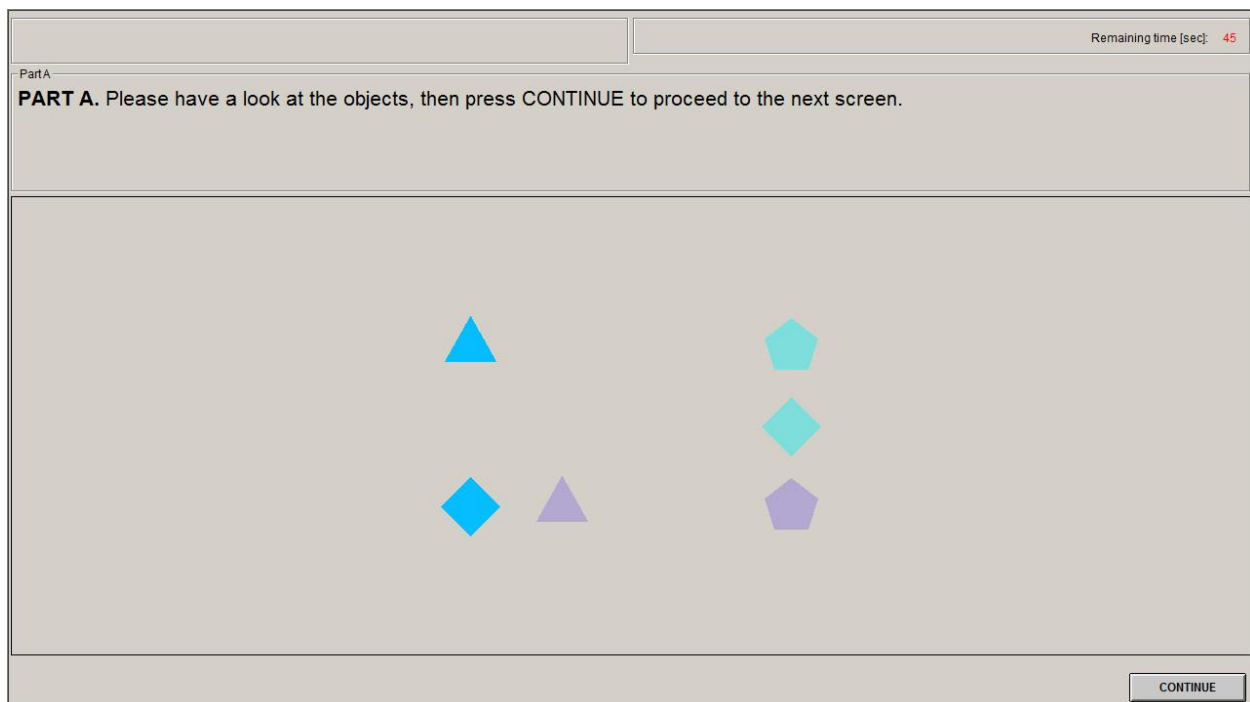
→Screenshot 3:



→Screenshot 4:

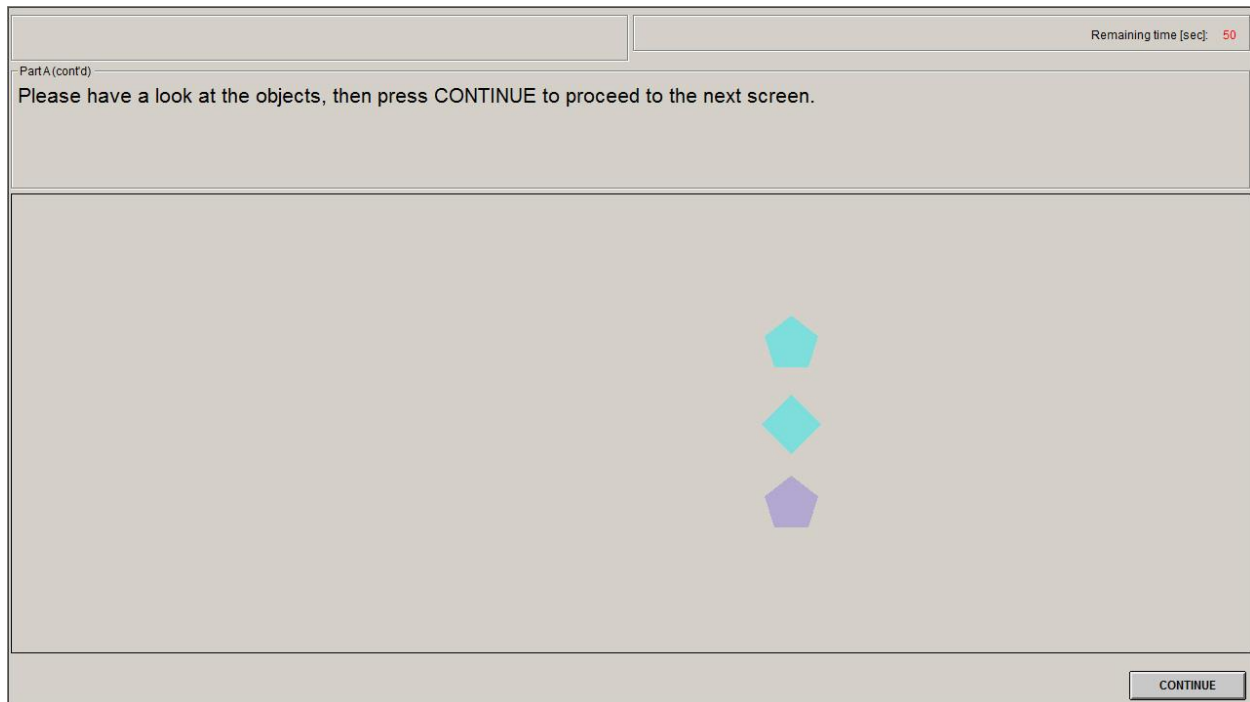


→Screenshot 5:



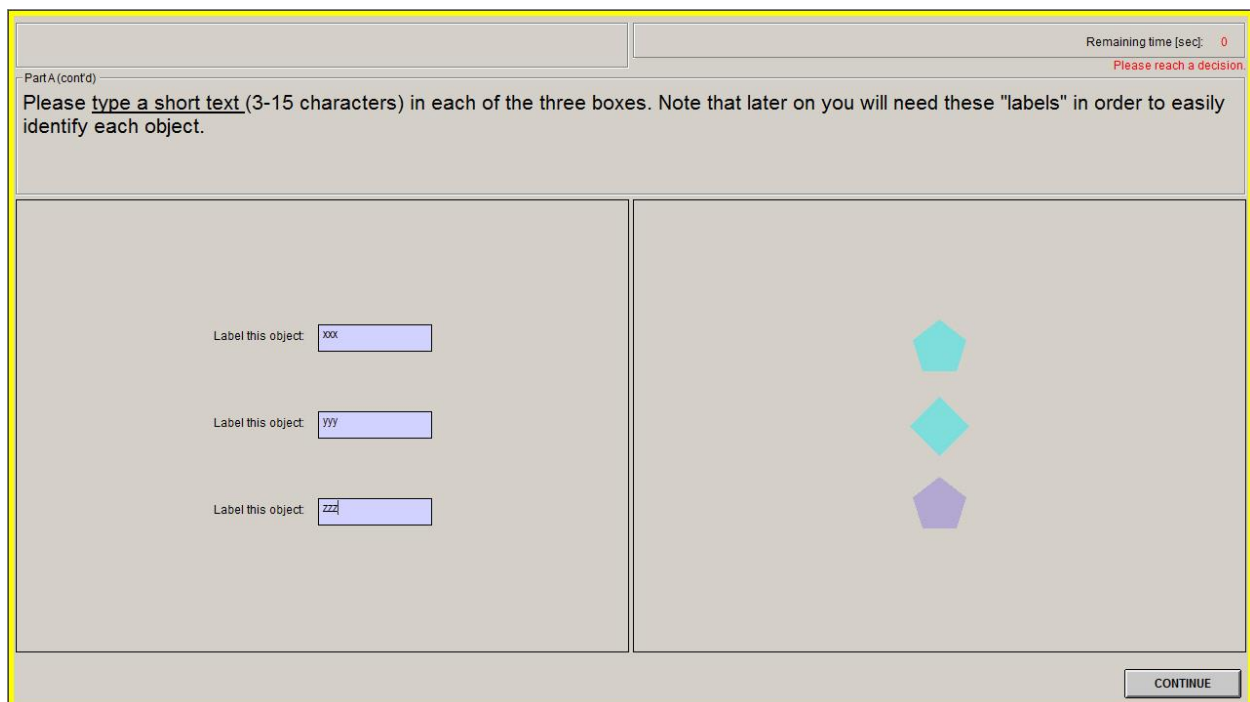
NOTE: Here three objects just slid toward the right; i.e., the program has selected three blocks – one by one – and put them in a column on the right-hand side of the screen.

→Screenshot 6:



NOTE: The (three) available strategy options are now shown on their own.

→Screenshot 7:



NOTE: The subject is prompted to enter one label per object, for her own future reference.

→Screenshot 8.i (*All-Aware* ONLY):

Part B

Remaining time [sec]: 20

Recall that – in Part C of the experiment – you will be prompted to pick one object in order to coordinate with your partner. Now, prior to that we would like to know what you think of the other participants in this room. Please answer the following questions by moving the sliders to the desired percentages. Note that your partner will not be asked to answer these questions.

1) How likely do you think it is that the other participants have noticed the different colors of the objects?
Please move the below slider.

50 %

2) How likely do you think it is that the other participants have noticed the different shapes of the objects?
Please move the below slider.

50 %

3) How likely do you think it is that the other participants have noticed the order in which the objects have been drawn by the computer program? Please move the below slider.

50 %

→Screenshot 8.ii:

Part B

Remaining time [sec]: 0
Please reach a decision.

At this point, we ask you to estimate how likely you think it is that the participants in this room (excluding yourself) will choose each of the objects. To express your estimates, please move the spokes to the desired percentages, and then press CONTINUE to proceed to the next screen.

Note: if at least one of your estimates differs by no more than 5 percentage points from the realized value, at the end of the study you will receive an additional payment of \$0.25.

Click and drag the small circles to move the spokes.

Object	Percentage
xxx	33%
zzz	33%
yyy	34%

NOTE: The spokes are initially arranged in such a way that each sector corresponds to one-third of the area. Within each sector there is a text box containing the label the subject entered in Screenshot 7.

→Screenshot 9:

Part C

PART C. Please pick one object. (Your partner has been shown the same set of objects as you.) Each of you and your partner will receive \$1.25 if you both pick the same object, \$0 otherwise.

You may pick your chosen object by ticking the box located next to the label.

Select your choice: xxx
 yy
 zzz

CONTINUE

NOTE: The labels next to each of the options are those entered by the subject in Screenshot 7.

→Screenshot 10 (transcript):

In what follows you will go through 9 other rounds, where each round will involve the same type of tasks as the ones you have performed so far.

NOTE: Screenshots 4 to 9 are repeated.